DataDirect NETWORKS

Storage Fusion Architecture

SFA10000/10000E
(SFA OS Version 1.4.0)

User Guide

96-00259-001 Rev C

Information in this document is subject to change without notice and does not represent a commitment on the part of DataDirect Networks, Inc. No part of this manual may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, for any purpose other than the purchaser's personal use without the written permission of DataDirect Networks, Inc.

© 2011 DataDirect Networks, Inc. All rights reserved.

DataDirect Networks, the DataDirect Networks logo, Silicon Storage Architecture, S2A, Web Object Scaler, and Storage Fusion Architecture are trademarks of DataDirect Networks, Inc. All other brand and product names are trademarks of their respective holders.

DataDirect Networks' Licensor(s) makes no warranties, express or implied, including without limitation the implied warranties of merchantability and fitness for a particular purpose, regarding the software. DataDirect Networks' Licensor(s) does not warrant, guarantee or make any representations regarding the use or the results of the use of the software in terms of its correctness, accuracy, reliability, currentness, or otherwise. The entire risk as to the results and performance of the software is assumed by you. The exclusion of implied warranties is not permitted by some jurisdictions. The above exclusion may not apply to you.

In no event will DataDirect Networks' Licensor(s), and their directors, officers, employees, or agents (collectively DataDirect Networks' Licensors) be liable to you for any consequential, incidental, or indirect damages (including damages for loss of business profits, business interruption, loss of business information, and the like) arising out of the use or inability to use the software even if DataDirect Networks' Licensor has been advised of the possibility of such damages. Because some jurisdictions do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitations may not apply to you. DataDirect Networks' Licensor's liability to you for actual damages from any cause whatsoever, and regardless of the form of the action (whether in contract, tort (including negligence), product liability or otherwise), will be limited to $50.

Document Number 96-00259-001 Rev C
Jun 2011

STANDARD WARRANTY

Definitions: This two-year limited warranty applies to the following DataDirect Networks network infrastructure and individual SAN solution components that include: Silicon Storage Appliance Hardware, Disk Modules, RAID Hardware Components, Storage Hardware Components, and Disk Docking Bays and Enclosures (hereinafter "DataDirect Networks Products"). Fibre Channel Interface Kits, SCSI Interface Kits, Host Adapters and Networking Products are limited to a 90-day warranty. Software bundled or included with DataDirect Networks solutions are furnished exclusively under the terms of the applicable license agreements.

Warranty: DataDirect Networks warrants that the DataDirect Networks Products accompanied by this limited Warranty are free from defects in material and workmanship for a period of two years from the date of original purchase from DataDirect Networks or an authorized DataDirect Networks reseller. During the term of this Warranty, DataDirect Networks will, at its option, repair or replace any defective parts of the DataDirect Networks products purchased under this Warranty at no additional charge. Repair parts or replacement DataDirect Networks products will be furnished on an exchange basis, and will be either reconditioned or new. When returning the DataDirect Networks products, the Purchaser must prepay any shipping charges. In addition, the Purchaser is responsible for insuring the products returned and assumes the risk of loss during shipment.

Warranty Claim Requirements: Purchaser claims made pursuant to this Warranty must conform to the following requirements:

1. The DataDirect Networks products must be returned to (a) an Authorized DataDirect Networks Servicing Reseller in the country of original purchase, or (b) a DataDirect Networks facility which performs Warranty service in the country of original purchase, or (c) an Authorized DataDirect Networks Third Party Service Provider in the country of original purchase.

2. The Purchaser must provide proof of purchase and date of purchase from DataDirect Networks or an Authorized DataDirect Networks Reseller.

3. The Purchaser may request information on how to obtain warranty service by contacting any Authorized DataDirect Networks Reseller, or by writing to the Warranty Service Department, DataDirect Networks, 9351 Deering Avenue, Chatsworth, CA 91311.

Disclaimers: THIS LIMITED WARRANTY DOES NOT APPLY TO ANY DATADIRECT NETWORKS PRODUCTS WHICH HAVE BEEN DAMAGED OR RENDERED DEFECTIVE (a) AS A RESULT OF ACCIDENT, MISUSE, OR ABUSE; (b) BY THE USE OF PARTS NOT MANUFACTURED OR SOLD BY DATADIRECT NETWORKS; (c) BY MODIFICATION WITHOUT THE WRITTEN PERMISSION OF DATADIRECT NETWORKS, OR (d) AS A RESULT OF SERVICE BY ANYONE OTHER THAN DATADIRECT NETWORKS, AN AUTHORIZED DATADIRECT NETWORKS SERVICING RESELLER, OR AN AUTHORIZED DATADIRECT NETWORKS THIRD PARTY SERVICE PROVIDER.

EXCEPT AS EXPRESSLY SET FORTH ABOVE, DATADIRECT NETWORKS MAKES NO OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, ANY IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR PURPOSE, AND DATADIRECT NETWORKS EXPRESSLY DISCLAIMS ALL WARRANTIES NOT STATED HEREIN. IN THE EVENT THE PRODUCTS ARE NOT FREE FROM DEFECTS AS WARRANTED ABOVE, THE PURCHASER'S SOLE REMEDY SHALL BE REPAIR OR REPLACEMENT AS PROVIDED ABOVE. UNDER NO CIRCUMSTANCES WILL DATADIRECT NETWORKS BE LIABLE TO THE PURCHASER, OR TO ANY USER, FOR ANY DAMAGES, EXPENSES, LOST PROFITS, LOST SAVINGS, DAMAGE TO OR REPLACEMENT OF EQUIPMENT AND PROPERTY, COSTS OF RECOVERING, REPROGRAMMING, OR REPRODUCING ANY PROGRAM OR DATA STORED IN OR USED WITH THE PRODUCTS, OR OTHER DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE DATADIRECT NETWORKS PRODUCTS.

ANY IMPLIED WARRANTIES ARE LIMITED TO THE TERMS OF THIS EXPRESS LIMITED WARRANTY. SOME STATES DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES FOR CONSUMER PRODUCTS, AND SOME STATES DO NOT ALLOW LIMITATIONS ON HOW LONG AN IMPLIED WARRANTY LASTS, SO THE ABOVE LIMITATIONS OR EXCLUSIONS MAY NOT APPLY TO YOU. THIS WARRANTY GIVES YOU SPECIFIC LEGAL RIGHTS, AND YOU MAY ALSO HAVE OTHER RIGHTS WHICH VARY FROM STATE TO STATE.

GNU GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/> Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program--to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps:

(1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS

0. Definitions.

"This License" refers to version 3 of the GNU General Public License.

"Copyright" also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

"The Program" refers to any copyrightable work licensed under this License. Each licensee is addressed as "you". "Licensees" and "recipients" may be individuals or organizations.

To "modify" a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a "modified version" of the earlier work or a work "based on" the earlier work.

A "covered work" means either the unmodified Program or a work based on the Program.

To "propagate" a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To "convey" a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays "Appropriate Legal Notices" to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

1. Source Code.

The "source code" for a work means the preferred form of the work for making modifications to it. "Object code" means any non-source form of a work.

A "Standard Interface" means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The "System Libraries" of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A "Major Component", in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The "Corresponding Source" for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

3. Protecting Users' Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any

non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

a) The work must carry prominent notices stating that you modified it, and giving a relevant date.

b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".

c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.

d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

a) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.

b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.

c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.

d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.

e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A "User Product" is either (1) a "consumer product", which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, "normally used" refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

"Installation Information" for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

7. Additional Terms.

"Additional permissions" are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or

b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or

c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or

d) Limiting the use for publicity purposes of names of licensors or authors of the material; or

e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or

f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered "further restrictions" within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An "entity transaction" is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

11. Patents.

A "contributor" is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's "contributor version".

A contributor's "essential patent claims" are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, "control" includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a "patent license" is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To "grant" such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. "Knowingly relying" means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify

or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is "discriminatory" if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others' Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found. <one line to give the program's name and a brief idea of what it does.> Copyright (C) <year> <name of author>

This program is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program. If not, see <http://www.gnu.org/licenses/>. Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode: <program> Copyright (C) <year> <name of author> This program comes with ABSOLUTELY NO WARRANTY; for details type `show w'. This is free software, and you are welcome to redistribute it under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, your program's commands might be different; for a GUI interface, you would use an "about box".

You should also get your employer (if you work as a programmer) or school, if any, to sign a "copyright disclaimer" for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see <http://www.gnu.org/licenses/>.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read <http://www.gnu.org/philosophy/why-not-lgpl.html>.

Copyright (C) 1997 Massachusetts Institute of Technology

This software is being provided by the copyright holders under the following license. By obtaining, using and/or copying this software, you agree that you have read, understood, and will comply with the following terms and conditions:

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose and without fee or royalty is hereby granted, provided that the full text of this NOTICE appears on ALL copies of the software and documentation or portions thereof, including modifications, that you make.

THIS SOFTWARE IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS. COPYRIGHT HOLDERS WILL BEAR NO LIABILITY FOR ANY USE OF THIS SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders. See the file AUTHORS which should have accompanied this software for a list of all copyright holders.

This file may be derived from previously copyrighted software. This copyright applies only to those changes made by the copyright holders listed in the AUTHORS file. The rest of this file is covered by the copyright notices, if any, listed below.

Copyright (c) 1992-1996 The Regents of the University of California. All rights reserved.

Permission to use, copy, modify, and distribute this software and its documentation for any purpose, without fee, and without written agreement is hereby granted, provided that the above copyright notice and the following two paragraphs appear in all copies of this software.

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS DOCUMENTATION, EVEN IF THE UNIVERSITY OF CALIFORNIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE UNIVERSITY OF CALIFORNIA SPECIFICALLY DISCLAIMS ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE PROVIDED HEREUNDER IS ON AN "AS IS" BASIS, AND THE UNIVERSITY OF CALIFORNIA HAS NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

Copyright (C) 2004, 2005, 2006 Board of Trustees, Leland Stanford Jr. University. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Preface

What is in this guide

This user guide contains comprehensive information regarding features and functions of the DataDirect Networks SFA10000/10000E system. It also gives you step-by-step instructions on how to install and configure the system. Information given in this user guide applies to both the SFA10000 and SFA10000E systems, unless otherwise stated.

Related Documentation

- SFA OS CLUI Command Reference
- SFA10000 Quick Start Guide
- SFA10000 Release Notes
- StorageScaler 6000 User Guide
- StorageScaler 7000 User Guide
- StorageScaler 2460 User Guide
- 42U/45U Rack Installation and Service Guide
- SFA10000 White Paper
- SFA OS Service Manual

International Standards

The SFA10000 complies with the requirements of the following agencies and standards:

- CE
- UL
- CUL
- C-Tick
- FCC

Potential for Radio Frequency Interference

USA Federal Communications Commission (FCC)

This equipment has been tested and found to comply with the limits for a class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. The supplier is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must

accept any interference received, including interference that may cause undesired operation.

European Regulations

This equipment complies with European Regulations EN 55022 Class A: Limits and Methods of Measurement of Radio Disturbance Characteristics of Information Technology Equipments and EN50082-1: Generic Immunity.

Canadian Regulations

ICES-003 Class A Notice - Avis NMB-003, Classe A

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

Safe Handling

- Remove drives to minimize weight
- Do not try to lift the enclosure by yourself
- Do not lift the SFA10000 by the handles at the front and on the power supply modules on the back; they are not designed to support the weight of the enclosure.

Safety

NOTE : If this equipment is used in a manner not specified by the manufacturer, the protection provided by the equipment may be impaired.

CAUTION ! Safety goggles should be worn when maintaining the equipment.

The SFA10000 *MUST* be grounded before applying power. Unplug the unit if you think that it has become damaged in any way and before you move it.

CAUTION ! To maintain proper airflow through the system, operate the system with the system top covers closed.

- Plug-in modules are part of the enclosure and must only be removed when a replacement can be immediately installed. The system must not be run without all modules in place.
- In order to comply with applicable safety, emission, and thermal requirements, the top covers should remain closed while running.
- The SFA10000 system must only be operated from a power supply input voltage range of 200VAC to 240VAC.

- A faulty power supply or fan module must be replaced with a fully operational module within 24 hours.

To minimize the risk of electric shock, disconnect the power from the power supply, either by turning off the switch or by physically removing the power cable, prior to removing the module from the enclosure.

- Do not remove a faulty power supply or fan module unless you have a replacement module of the correct type ready for insertion.
- The power connection must always be disconnected prior to removal of the power supply module from the SFA10000 or disk enclosures.
- A safe electrical earth connection must be provided to the power cord.
- Provide a suitable power source with electrical overload protection to meet the requirements given in the technical specifications.

Do not remove covers from the power supply module. Danger of electric shock inside. Return the module to your supplier for repair.

Operation of the SFA10000 with *ANY* modules missing will disrupt the airflow and the components will not receive sufficient cooling. It is *ESSENTIAL* that all apertures are filled before operating the unit.

Recycling of Waste Electrical and Electronic Equipment (WEEE)

At the end of the product's life, all scrap/ waste electrical and electronic equipment should be recycled in accordance with National regulations applicable to the handling of hazardous/ toxic electrical and electronic waste materials.

NOTE : Observe all applicable safety precautions, such as weight restrictions, handling batteries and lasers etc, detailed in the preceding paragraphs when dismantling and disposing of this equipment.

ESD Precautions

CAUTION ! It is recommended that you check and fit a suitable anti-static wrist or ankle strap and observe all conventional ESD precautions when handling the SFA10000 plug-in modules and components. Avoid contact with backplane components and module connectors.

Data Security

- Disk units are fragile. Handle them with care, and keep them away from strong magnetic fields.
- ***ALL*** the supplied plug-in modules and blanking plates must be in place for the air to flow correctly around the enclosure and also to complete the internal circuitry.
- If the SFA10000 or disk enclosure is used with modules or dummy disk modules missing for more than a few minutes, the system can overheat, causing power failure and data loss. Such use may also invalidate the warranty.
- If you remove a drive module, replace it immediately. If it is faulty, replace it with a drive module of the same type and capacity.
- Ensure that all disk drives are removed from the enclosure before attempting to move the rack installation.
- Do not abandon your backup routines. No system is completely foolproof.

Table of Contents

Preface

Related Documentation xii
International Standards xii
Potential for Radio Frequency Interference xii
European Regulations xiii
Canadian Regulations xiii
Safe Handling xiii
Safety xiii
Recycling of Waste Electrical and Electronic Equipment (WEEE) xiv
ESD Precautions xiv
Data Security xv

Chapter 1
Introduction

1.1 Introduction 2
1.1.1 DataDirect Networks' Storage Fusion Architecture (SFA) 2
1.1.2 Product Variations 2
1.1.3 Features of the SFA10000 3

1.2 The SFA10000 System Hardware 5
1.2.1 Status LED Indicators 7
1.2.2 Fan Module 8
1.2.3 Power Supply Module 8
1.2.4 Internal Disk Modules 9
1.2.5 RAID Processor and Application Processor 9
1.2.6 I/O Ports 10
1.2.7 Uninterruptible Power Supplies (UPS) 12

Chapter 2
Installation

2.1 Installation Overview 14

2.2 Site Preparation 15
2.2.1 Delivery Route Verification 15
2.2.2 Rack Location, Air Flow, and Access 16
2.2.3 Floor Loading 17
2.2.4 Cooling Supply Planning 17
2.2.5 AC Power Supply Planning and Verification 17

2.3 Unpacking the SFA10000 19
2.3.1 Packing List 19

2.4 Installing the Disk Modules 20

2.5 Cable Connections 21
2.5.1 Couplet ICL Cabling 21
2.5.2 Disk Enclosure Cabling 22
2.5.3 Host Connections 29
2.5.3.1 SFA10000 29

2.5.3.2 SFA10000E 29
2.5.4 Management Network Connection 30
2.5.5 RS-232 Console Connection 30
2.5.6 UPS Connection 31
2.5.6.1 UPS Battery 31
2.5.7 Power Connections 32

2.6 Powering On the System 32

2.7 Configuring the SFA10000 34
2.7.1 Planning Your Setup and Configuration 34
2.7.2 Serial Interface Configuration 35
2.7.3 Validate the Hardware 35
2.7.4 Clear System Configuration (SFA10000 only) 37
2.7.5 Set the System Name 37
2.7.6 Set System Time & Date (NTP) 38
2.7.7 Configure Network Interface Settings 38
2.7.8 Access Virtual Machines (SFA10000E Only) 39
2.7.9 Create Storage Pools 40
2.7.10 Create Virtual Disks 42
2.7.11 Create and Assign Spare Pools 44
2.7.12 Present Virtual Disk to External Host (SFA10000 Only) 46
2.7.12.1 Special Considerations for MAC OS® 48
2.7.12.2 Selective Presentation 49
2.7.13 Storage Pool Initialization 49

Chapter 3
Administration

3.1 Managing the SFA10000 51
3.1.1 Management Interface 51
3.1.2 User Logins 52
3.1.3 Available Commands 52
3.1.3.1 Basic Key Operations 53

3.2 Physical Disk Information 54
3.2.1 Disk Information 54
3.2.2 Disk States 55
3.2.3 Visual Indication 56

3.3 Storage Pool Management 57
3.3.1 Display Storage Pool Information 57
3.3.2 Creating a Storage Pool 59
3.3.3 Storage Pool Initialization 60
3.3.3.1 Initialization Job Failure 60
3.3.4 SATAssure 61
3.3.5 Verifying a Storage Pool 62
3.3.6 Naming a Storage Pool 62
3.3.7 Deleting a Storage Pool 62
3.3.8 Locate a Storage Pool 62

3.4 Virtual Disk Management 63
3.4.1 Display Virtual Disk Information 63
3.4.2 Creating a Virtual Disk 64
3.4.3 Naming a Virtual Disk 64
3.4.4 Deleting a Virtual Disk 64
3.4.5 Presentation to Hosts (not applicable on SFA10000E) 64

3.5 Spare Pool Management 65
3.5.1 Display Spare Pool Information 65
3.5.2 Creating a Spare Pool 66
3.5.3 Naming a Spare Pool 66
3.5.4 Deleting a Spare Pool 66

3.6 Presentations 67
3.6.1 Discovered Initiator Commands 68
3.6.2 Host Commands 68
3.6.3 Presentation Commands 68
3.6.3.1 Persistent Reservation Support 69
3.6.4 Additional Configuration Considerations for Mac® Hosts 69
3.6.5 Configure Presentations of Virtual Disks to Hosts 70
3.6.5.1 Create Host Object 71
3.6.5.2 Identifying Host FC Connections via Ports 72
3.6.5.3 Import Discovered Initiators 75
3.6.5.4 Present a Virtual Disk to a Host 75

3.7 Network Time Protocol Mode 78

3.8 Disk Rebuild 80
3.8.1 Full and Partial Rebuilds 80
3.8.2 Sparing Policy 80
3.8.3 Manual Disk Replace/Rebuild 81
3.8.4 Manual Fail/Rebuild of a Disk 81

3.9 Restarting the SFA10000 82
3.9.1 System Restart 82
3.9.2 System Shutdown 82

3.10 Performance Management 83
3.10.1 Cache Coherency 83
3.10.1.1 Cache Protection 83
3.10.2 Single Controller Write Back Cache Policy 83
3.10.3 Right Side I/O 84
3.10.4 Background Job Priority 85
3.10.4.1 Pause/Resume a Job 86
3.10.5 Rebuild Policy Priority 86

3.11 Firmware Update Management 88
3.11.1 Displaying Current Firmware Version 89
3.11.2 Controller Firmware Update Procedure 89

3.11.2.1 Linux Environment-Firmware Update from the Network 90
3.11.2.2 Windows Environment-Firmware Update from the Network 90
3.11.3 Disk Enclosure Firmware Update Procedure 92
3.11.3.1 Enclosure File Definitions 92
3.11.3.2 Linux Environment-Firmware Upload 93
3.11.3.3 Windows Environment-Firmware Upload 93
3.11.3.4 Disk Shelf Upgrade 93
3.11.4 Physical Disk Firmware Update Procedure 96
3.11.4.1 Linux Environment-Firmware Upload 96
3.11.4.2 Windows Environment-Firmware Upload 97
3.11.4.3 Physical Disk Upgrade 97

3.12 The System Logs 99
3.12.1 Displaying Event Logs 99
3.12.2 Event Log Structure 100
3.12.3 SFA OS Terms in the Log 100

3.13 Remote Management of SFA10000 101
3.13.1 Network Connection 101
3.13.2 Display Network Interface Attributes 101
3.13.3 Change Network Interface Settings 101
3.13.4 Logins 102
3.13.5 Email and SNMP Notification Setup 102
3.13.5.1 Email Setup 102
3.13.5.2 SNMP Setup 103
3.13.5.3 Inquiry Items and Events 103

Chapter 4
GUI Management Agent

4.1 Starting the GUI Management Agent 107
4.1.1 Login 107
4.1.2 Home Screen and Health Indicators 108

4.2 Subsystem Menu 109
4.2.1 Show Subsystem 109
4.2.2 Set Subsystem Attributes 110
4.2.2.1 NTP Settings 110
4.2.3 Show Background Jobs 110
4.2.4 Set Background Jobs Attributes 111
4.2.5 Email and Critical Event Notification Setup 111

4.3 Controllers Menu 113
4.3.1 Show Controllers 113
4.3.2 Set Controller Attributes 113
4.3.3 Update Controller Firmware 114
4.3.4 View System Logs 114

4.4 Physical Disks Menu 115
4.4.1 Show List of Physical Disks 115
4.4.2 Locate Disks 116

4.4.3 Remove Disk from Spare Pool 116
4.4.4 Assign Disk to Pools 116

4.5 Unassigned Pool 117
4.5.1 Show Unassigned Pool 117
4.5.2 Locate Unassigned Pool 118

4.6 Storage Pools 119
4.6.1 Show List of Storage Pools 119
4.6.2 Create a Storage Pool 120
4.6.3 Configure Storage Pool Attributes 121
4.6.4 Locate Storage Pools 122
4.6.5 Delete a Storage Pool 122

4.7 Virtual Disks Menu 123
4.7.1 Show List of Virtual Disks 123
4.7.2 Create a Virtual Disk 124
4.7.3 Configure Virtual Disk Attributes 124
4.7.4 Delete a Virtual Disk 125

4.8 Spare Pools 126
4.8.1 Show Spare Pools 126
4.8.2 Locate a Spare Pool 127
4.8.3 Create a Spare Pool 127
4.8.4 Assign Spare Pool to Storage Pool 128
4.8.5 Set Spare Pool Attributes 129
4.8.6 Delete a Spare Pool 129
4.8.7 Delete a Disk from Spare Pool 129

4.9 Presentations 130
4.9.1 Presentation Menu 131
4.9.2 Show Presentations 131
4.9.3 Set Up a Presentation 131
4.9.3.1 Create a Host 131
4.9.3.2 Map Host to Discovered Initiators 132
4.9.3.3 Present Virtual Disks to Hosts 132

4.10 Enclosure Menu 133
4.10.1 Show Enclosures 133
4.10.2 Update Enclosure Firmware 133

Chapter 5
Support

5.1 "RAID[0]$" Prompt is not shown correctly 136

5.2 Event Log 136

5.3 Manual Intervention Required (MIR) States 137

5.4 Recovery from Disk Failure 139
5.4.1 Automatic Rebuild 139

5.4.2 Spare Disk Replacement 140
5.4.3 When a Spare is not available 140
5.4.4 Manual Rebuild 140
5.4.5 Power Cycling Disk Drives 141

5.5 UPS Maintenance 142
5.5.1 UPS and Battery Pack Care 142
5.5.2 Checking the UPS Status 142
5.5.3 Replacing the Battery Pack 143
5.5.3.1 Testing Batteries 145

5.6 Controller Component Failures 146
5.6.1 Replacing a Power Supply Module 147
5.6.2 Replacing a Fan Module 148
5.6.3 Replacing an Internal Disk Module 149

5.7 Disk Enclosure Component Failures 150
5.7.1 I/O Module Replacement 150
5.7.2 DEM Replacement on SS6000 151
5.7.3 SAS Cable Replacement 152

Appendices

Appendix A. Technical Specifications 154

Appendix B. Disk Module Installation 157
B.1 SS6000 Disk Module Installation Procedure 157
B.2 SS7000 Disk Module Installation Procedure 158
B.3 SS2460 Disk Module Installation Procedure 159

Appendix C. SS6000 Disk Maps for SFA10000E 160

Appendix D. SS6000 Zoning Configuration Procedure 161

Appendix E. Rack Installation Procedure 163
E.1 Rackmounting the Controller 163
E.2 Rackmounting the UPS Unit 169

Appendix F. FRU Part Numbers 171

Appendix G. Terminology 173

Index 175

Contacting Technical Support & Shipping Instructions 178

Chapter 1

Introduction

1.1 Introduction

Welcome to the DataDirect Networks SFA10000 User Guide.

1.1.1 DataDirect Networks' Storage Fusion Architecture (SFA)

DDN's Storage Fusion Architecture provides the foundation for balanced, high-performance storage system. Using highly parallelized storage processing technology, SFA delivers unprecedented IOPS and massive throughput at the same time. Combined with Enterprise-level data protection and leading scalability, density and energy efficiency, SFA is the storage technology of choice for the most extreme storage infrastructure projects.

The SFA10000 employs RAID, data integrity and data management software written from the ground up to take advantage of multi-core processors and modern bus architectures. Its highly threaded architecture allows performance to linearly scale with advances in underlying hardware. This same architecture allows the SFA10000 to perform in the extreme range of both throughput and IOPS. Designed to house the most scalable unstructured file data, the system supports up to 1200 disks of raw storage while enabling a combination of SAS, SATA or SSD disks.

1.1.2 Product Variations

Various configurations (Figure 1) are available with the SFA10000 to enable full system performance, capacity, and the highest levels of disk enclosure fault tolerance.

- Five 60-bay disk enclosures with up to 300 disk modules at 28U in one rack
- Ten 60-bay disk enclosures with up to 600 disk modules at 48U in two racks
- Twenty 60-bay disk enclosures with up to 1200 disk modules at 88U in two racks
- Five 24-bay disk enclosures with up to 120 disk modules at 18U in one rack

Figure 1. SFA10000 Possible Configurations

1.1.3 Features of the SFA10000

The SFA10000 incorporates the following features:

- 8GB/s full-duplex cache link
- Internal SAS Switching (480Gb/s Internal SAS Storage Network)
- Up to 1200 SAS, SATA, or SSD disks with full redundant paths
- **InfiniBand™ (IB) or Fibre Channel (FC) Connectivity on SFA10000**
  For SFA10000 IB option, the two Controllers provide up to eight (8) QDR InfiniBand host port connections.
  For SFA10000 FC option, the two Controllers provide up to sixteen (16) individual 8Gb/s Fibre Channel host port connections, including simultaneous access to the same data through multiple ports. Each FC host port supports point-to-point and switched fabric operation.
- **InfiniBand or 10G Ethernet Connectivity on SFA10000E**
  The SFA10000E provides up to 16 individual IB or10Ge ports for host access to the Virtual Machines. Each port can be configured to run as either InfiniBand QDR or 10G Ethernet. 10G Ethernet connectivity requires special cabling that can be purchased separately from DDN. See your sales contact for more details.
- **Active/Active Operation with Cache Coherency and Failover**
  SFA implements an active/active host presentation model with routing-based data access and full cache coherency.
- **Data Protection**
  The SFA RAID stack provides protection against single physical disk failures with RAID 1 or RAID 5 data protection as well as double physical disk failures through the use of high-speed RAID 6 protection.
- **Configurable RAID Group Sizes**
  (5 or 9 disks/RAID 5 group, 6 or 10 disks/RAID 6 group, 2 disks/RAID 1 group)
  This feature allows you to configure the system with the desired RAID and redundancy levels based on the importance of your data. Each RAID group is configured independently and any valid combination for the number of disks in the array is supported.
- **SAS/SATA/SSD storage pool**
  Intermixing high-performance SSD and SAS disks with high-capacity SATA disks provides the flexibility to migrate disk technology along with varying application needs. The SFA10000's enclosure-level disk intermixing and exceptional per-disk performance reduces the number of disks needed to meet growing performance and capacity demands—lowering both acquisition and operation costs.
- **SATAssure™ Data Protection**
  SATAssure technology (Silent Data Corruption Detection and Avoidance) is designed by DDN to improve the reliability of enterprise SATA disks and make sure that data integrity is always mentioned for all I/O operations.
- **Battery Backed Write Back Cache**
  SFA OS provides a write back cache feature that is used to improve I/O performance. Write back cache data that has not been written to disk is preserved by maintaining power to the cache memory in the event of an AC mains failure long enough to copy the contents of the cache to stable storage. In addition, SFA OS is designed to tolerate a simultaneous AC mains failure and RAID software failure.

- **Mirrored Write Back Cache**
  Currently, SFA OS provides the ability to mirror all write back cache data such that the failure of a single controller will not result in data loss. A storage administrator can optionally turn off write back cache mirroring for a RAID set (for higher performance); however data protection is reduced for logical units within that RAID set.

- **Mirrored Transaction Journal**
  RAID write holes are prevented by executing stripe updates as ACID (Atomicity, Consistency, Isolation, Durability) transactions. If the transactions are interrupted by a power failure, they can be recovered from the transaction journal implemented within the write back cache when power is restored. This journal is mirrored so that when a simultaneous power failure and controller hardware failure occurs, the surviving controller can recover the transactions.

- **Storage Metadata Mirrored n-Ways**
  SFA OS stores a copy of storage system metadata on 18 physical disks to minimize the likelihood that its metadata is lost or corrupted.

- **Partial Rebuild**
  This feature reduces rebuild times by updating only the data that has changed while the disk was down. The SFA OS tracks the changes made to a RAID set when a member physical disk becomes unavailable. If that member becomes available again within a user-definable timeout, then only the stripes that were modified while the member was missing are rebuilt. This minimizes the mean-time-to-repair for the RAID set and thus limits any performance impact of a disk repair.

- **Hot-Swappable and Redundant Components**
  Adhering to enterprise RAS standard, almost all hardware components (such as physical disks, power supply modules, and fan module) are redundant and hot-swappable.

- **Hot Spares**
  The SFA OS provides pools of spare physical disks that can be automatically used to replace failed physical disks. By replacing a failed RAID set member automatically, the mean-time-to-repair for the RAID set is minimized resulting in improved data reliability.

- **Management Options via RS-232 and Ethernet (SSH)**
  A RS-232 port and Ethernet ports are included to provide local and remote management capabilities.

1.2 The SFA10000 System Hardware

This section describes the hardware components of the SFA10000 controller. Refer to *DDN SS6000 User Guide*, *DDN SS7000 User Guide*, and *DDN SS2460 User Guide* for hardware descriptions on the SS6000, SS7000, and SS2460 disk enclosures.

The SFA10000 controller is a three-unit (3U), rack-mountable enclosure (Figure 2).

Figure 2. SFA10000 Controller Chassis

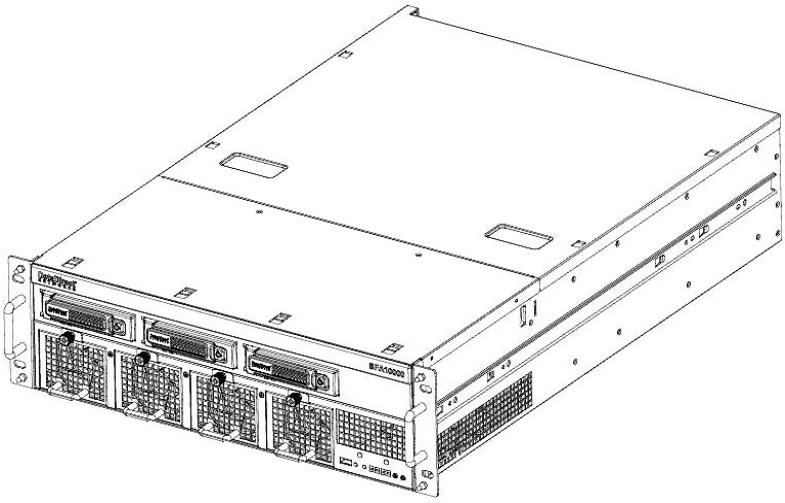

At the front, there are four fan modules, three internal disk modules, control buttons and status LED indicators (Figure 3).

Figure 3. SFA10000 Controller Front View without Bezel

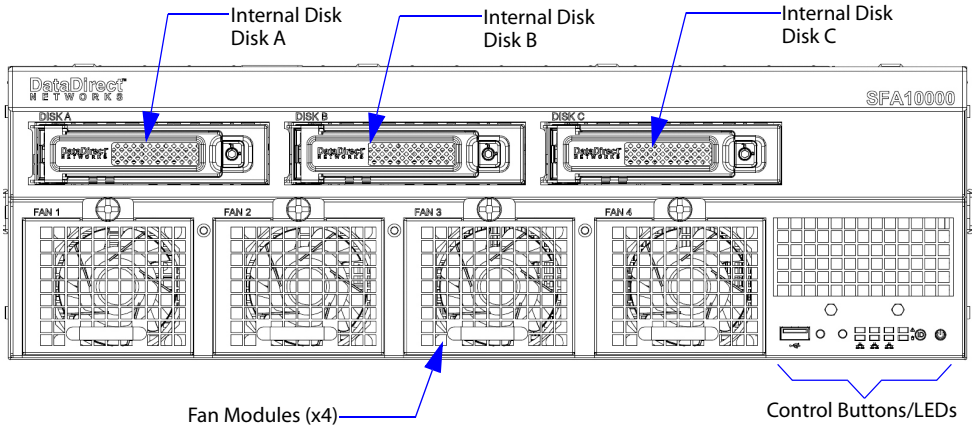

At the back, there are two power supply modules and various I/O connectors. Figure 4 and Figure 5 illustrate the rear views of the SFA10000 and SFA10000E controllers respectively.

Figure 4. SFA10000 Controller Rear View

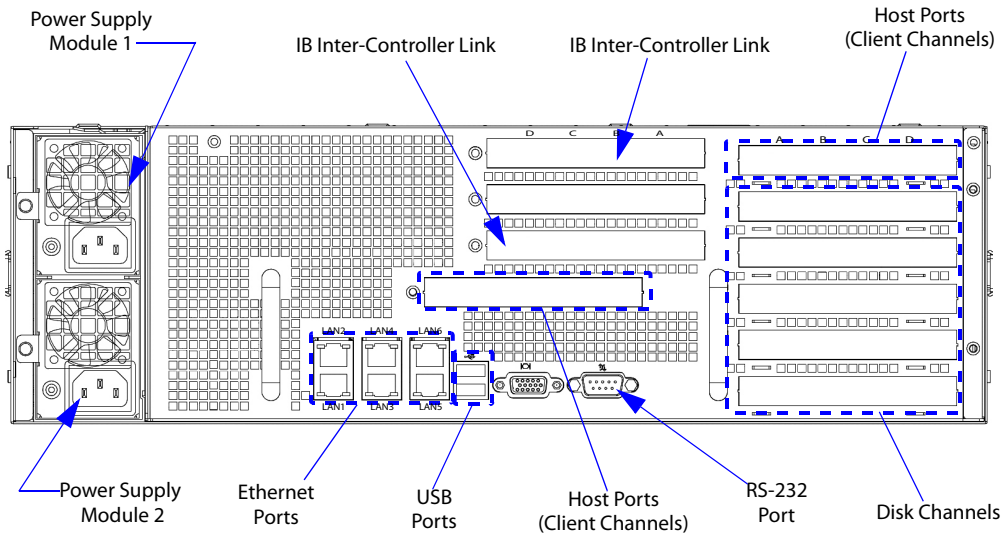

Figure 5. SFA10000E Controller Rear View

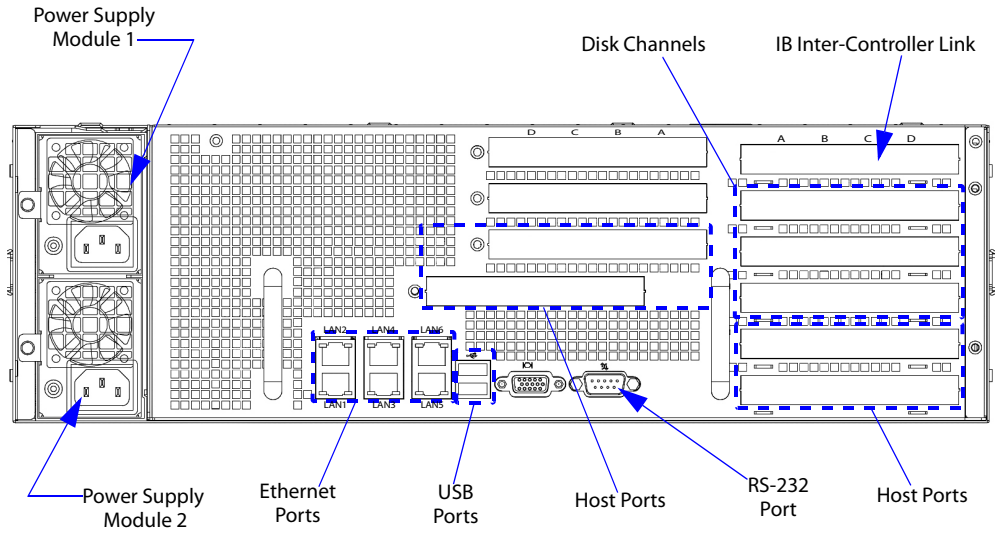

1.2.1 Status LED Indicators

Figure 6 below illustrates the positions of all the control buttons and LED indicators on the front panel. The LEDs are defined in Table 1. The USB port can be used for downloading new firmware and BIOS to the controller.

Figure 6. Controller Front Panel Control Buttons and Status LED Indicators

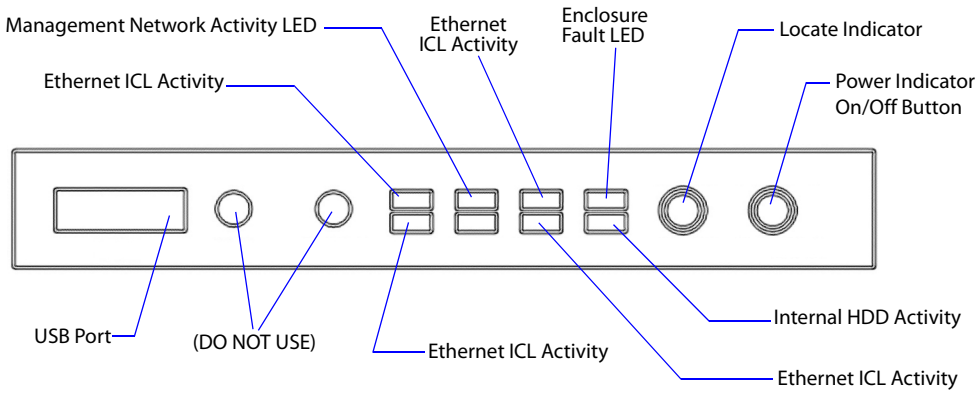

NOTE : Do **NOT** use the Power button under normal operation since doing so may cause data loss. If you need to power down the controller, issue the SHUTDOWN command via CLUI or GUI first, then use the UPS's Power button to turn off power to the controller. Refer to Section 3.9.2 "System Shutdown" for more information.

Table 1. Controller Status LED Indicators

Description	Color	Status
Power	Green	ON - power is applied to controller enclosure OFF - no power is applied to controller enclosure
Locate	Blue	Flashing at a 2-second interval - receiving "Locate Enclosure" command OFF - NOT receiving "Locate Enclosure" command
Enclosure Fault	Amber	ON - a fan failure, power supply failure, or over temperature condition occurred. A service action is required. OFF - no detectable faults
HDD	Green	Flashes - indicates HDD activity OFF - no HDD activity
Ethernet ICL	Green	Flashes - indicates the corresponding ICL link is up and both controllers are running OFF - no ICL activity
Network Activity	Green	ON - good connection is established on the corresponding Ethernet port Flashes - indicates activity

1.2.2 Fan Module

Each controller is equipped with four fan modules (Figure 7). These fan modules provide redundant cooling system for the unit. If one module fails, the other three will maintain sufficient cooling for the enclosure. The faulty module will still be providing proper air flow for the enclosure so do not remove it until a new module is available for replacement.

Figure 7. Controller Fan Module

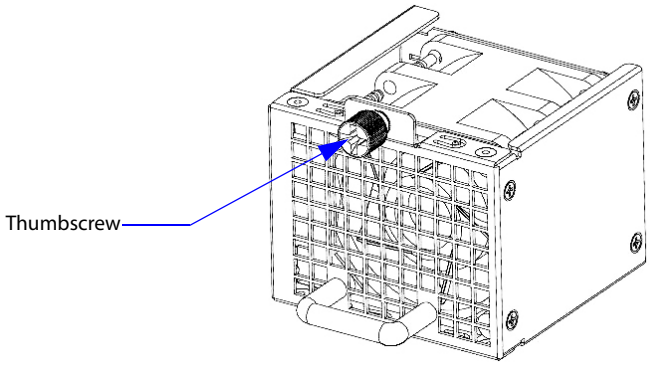

1.2.3 Power Supply Module

Each controller includes two power supply modules (Figure 8). These modules are redundant and hot-swappable. If one module fails, the other module will maintain sufficient power to the enclosure. The faulty module will still be providing proper air flow for the enclosure so do not remove it until a new module is available for replacement.

The LED mounted on the module indicates the status of the power supply. It is green when the module is operating normally and turns off when a fault occurs.

Figure 8. Controller Power Supply Module

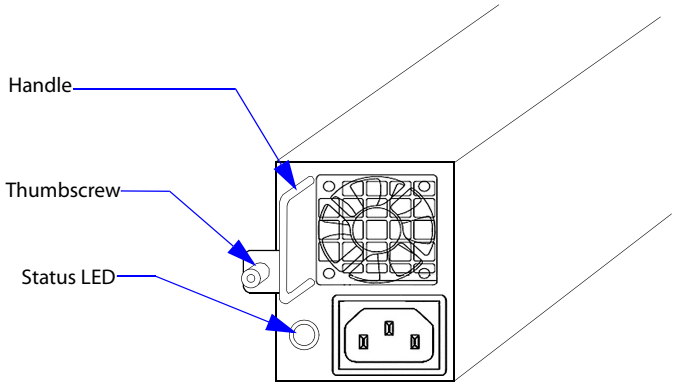

1.2.4 Internal Disk Modules

The SFA10000 controller includes three internal hard disks. Disk A and Disk B are mirrored system disks to run the controller. The information on these 2 disks includes the Linux OS, the SFA OS software, the event log, and controller-specific configuration (not storage configuration). Disk C is used for a large trace log[a].

Each disk module comprises a single low profile 1.0-inch high, 3.5-inch form factor hard disk mounted in a carrier. The module handle provides camming of the module into and out of disk bays and positive "spring loading" of the disk/baseplane connector (Figure 9). The handle is released by pressing the handle release button.

Each disk module has 2 LEDs—one green and one blue. Flashing green indicates disk activity while solid green indicates that the disk is online but idle. When SFA OS starts, it lights all 3 blue LEDs for 4 seconds. During discovery and while in a MIR state, the blue LEDs alternately blink from right to left. Under normal condition, the blue LEDs alternately blink from left to right. If all of the blue LEDs are off, then SFA OS is not completely running. Common failures will result in a single blue LED being on or all blue LEDs being off.

Figure 9. Controller Internal Disk Module

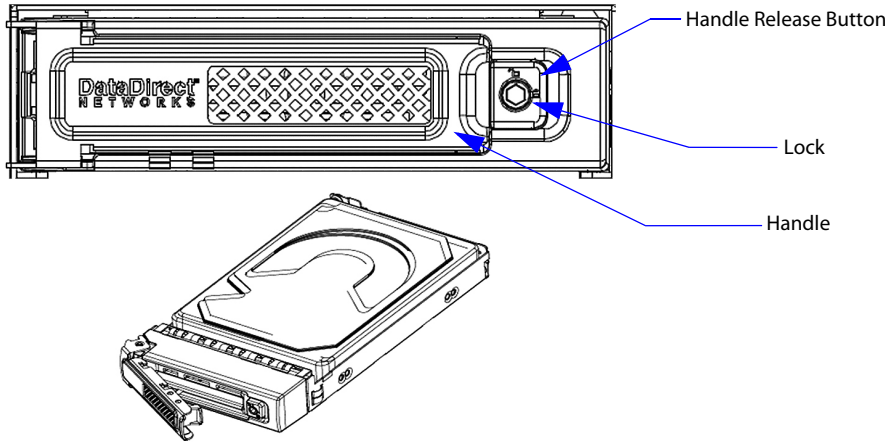

1.2.5 RAID Processor and Application Processor

The two controllers are redundant and hot-swappable which provide the intelligence and active/active data protection features of the SFA10000. If a controller fails, the remaining controller will assume its functionality and continue to provide data access, at a reduced performance level.

The SFA10000 has two parallel, multi-threaded RAID engines that work simultaneously in each controller for a total of 4 RAID Processors across the redundant controller pair.

The SFA10000E has a single, parallel, multi-threaded RAID engine in each controller, for a total of 2 RAID Processors across the redundant controller pair. In addition to the RAID Processor, each SFA10000E controller has a dedicated Application Processor. This processor is responsible for managing and running the Virtual Machines, with each Virtual Machine receiving one or more CPU cores.

a. The mirrored system disk feature is not yet implemented in V1.4.0. Currently, Disk A is the system disk and Disks B and C are not used.

1.2.6 I/O Ports

Figure 10 and Figure 11 below illustrate the I/O ports on the back of the SFA10000 and SFA10000E controllers respectively.

The InfiniBand ICL and Ethernet ICL ports provide connections for Mirrored Write Back Cache and control between the controllers in a couplet configuration.

The Disk I/O Channels are used for disk enclosure connections.

The UPS port is used for UPS connection.

The RS-232 connector provides local system monitoring and configuration capabilities. The VGA monitor and USB keyboard ports can be used as an alternative to the RS-232 console. However, the RS-232 console is recommended since its output can be logged and its connection can be longer.

Figure 10. SFA10000 Controller I/O Ports

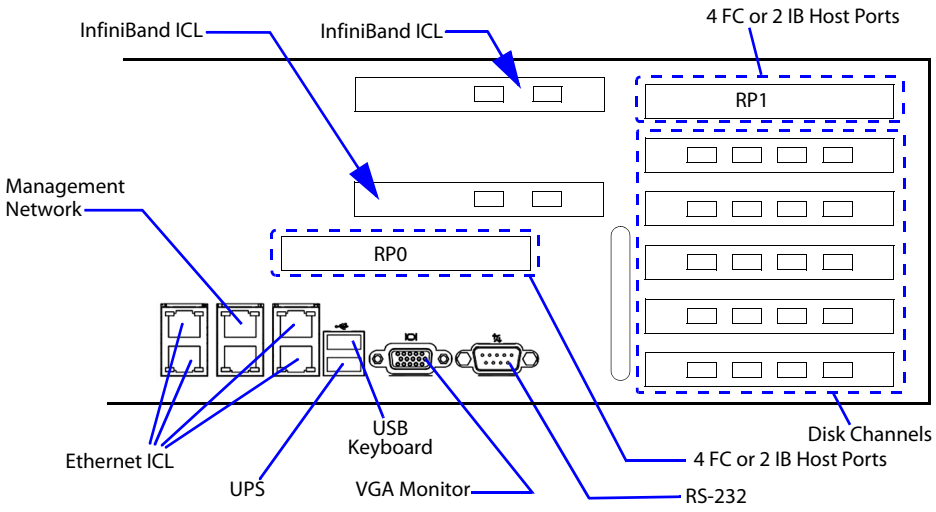

Figure 11. SFA10000E Controller I/O Ports

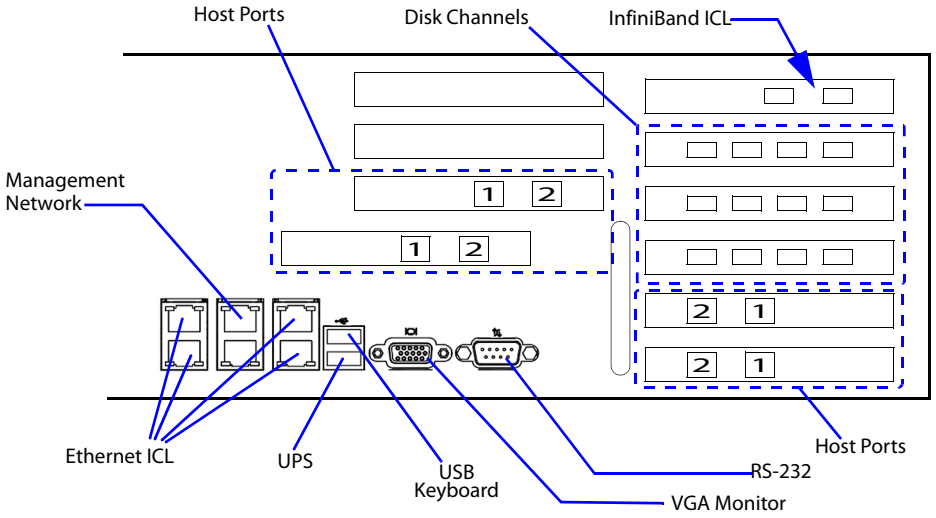

The Management Network Ethernet port provides remote monitoring and configuration capabilities. The colors of the LEDs indicate the link speed and activities (Figure 12).

Figure 12. Management Network Ethernet Port LED Color Scheme

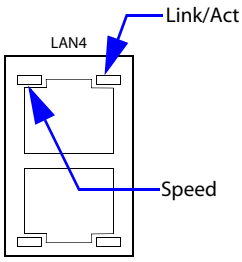

LED	Color	Status
Speed	Solid Orange	1000 Mb/s
	Solid Green	100 Mb/s
	Off	10 Mb/s
Link/Activity	Solid Green	Link established and idle
	Flashing Green	Link established and active
	Off	Link down

On the SFA10000, the RP0 and RP1 host ports (client channels) provide 8 Fibre Channel or 4 InfiniBand host connections on each controller. The LED color schemes are described in Figure 13 and Figure 14.

On the SFA10000E, there are 8 InfiniBand/10Ge host connections available on each controller. The IB host port LED color schemes are described in Figure 14.

Figure 13. SFA10000 FC Host Port LED Color Scheme

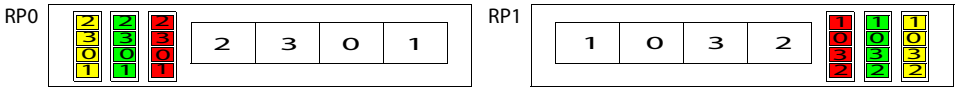

Yellow LED (8 Gbps)	Green LED (4 Gbps)	Red LED (2 Gbps)	Status
Off	Off	Off	Power off
On	On	On	Power on (before firmware init)
Flashing	Flashing	Flashing	Power on (after firmware init)
Yellow, Green, and Red LEDs flashing alternatively			Firmware error
Off	Off	On/Flashing	Online, 2 Gbps link / I/O activity
Off	On/Flashing	Off	Online, 4 Gbps link / I/O activity
On/Flashing	Off	Off	Online, 8 Gbps link / I/O activity
Flashing	Off	Flashing	Locate

Figure 14. SFA10000/10000E IB Host Port LED Color Scheme

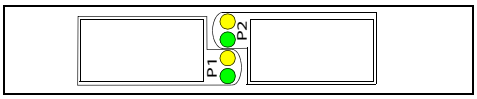

Description	Color	Status
Physical link	Green	ON - good physical link Flashing - indicates a problem with the link
Data activity	Yellow	ON - no data transfer Flashing - indicates data activity

1.2.7 Uninterruptible Power Supplies (UPS)

Each controller within the SFA10000 system is paired with a UPS to provide battery backup power when the AC mains fail. The UPS is a rack-mountable 1U unit.

Figure 15 below illustrates the positions of all the control buttons and LED indicators on the front panel. The DIP switches on the back panel are set at the factory and should only be changed by the DDN field engineer (Figure 16).

Figure 15. UPS Front View

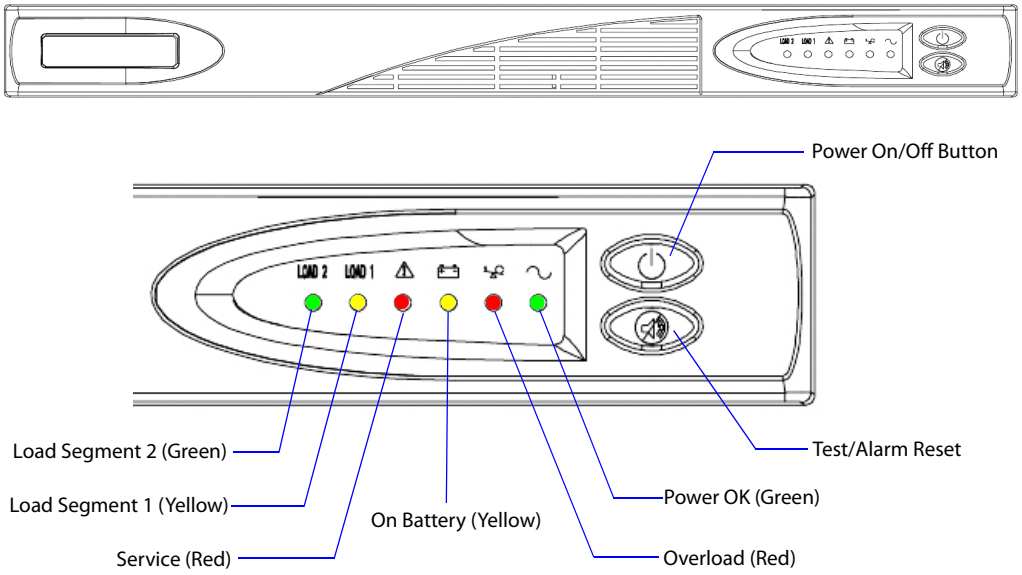

Figure 16. UPS Rear View

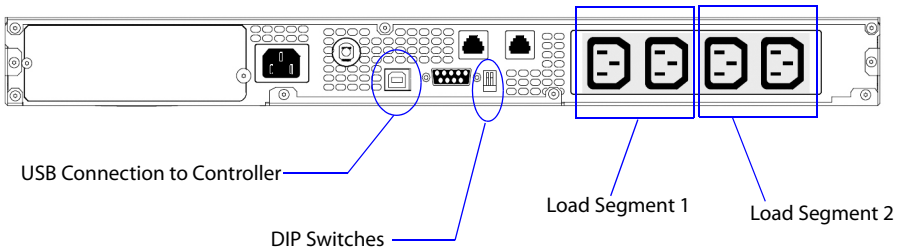

Chapter 2

Installation

2.1 Installation Overview

Here is an overview of all the steps needed to be taken to complete an installation and configuration of an SFA10000/10000E system.

> **NOTE :** All installation procedures apply to both the SFA10000 and SFA10000E systems, unless otherwise stated.

1. Site preparation.
2. Unpack the system.
3. Install the disk modules into the disk enclosures.
4. Connect the cables.
5. Power up the system.
6. Verify zoning configuration on disk enclosures.
7. Validate the hardware.
8. Configure the network interface settings via the RS-232 console.
9. For SFA10000, configure the storage (create storage pools, virtual disks, spare pools, and presentations).

   For SFA10000E, use either a VNC viewer or SSH to access the preconfigured Virtual Machines.

2.2 Site Preparation

This section provides information on how to prepare a site prior to installing a SFA10000 system. The site preparation steps include:

- Delivery route verification
- Rack location, air flow, and access
- Floor loading
- Cooling supply planning
- AC power supply planning and verification

2.2.1 Delivery Route Verification

Each SFA10000 system consists of one to two 19" racks. Each of these racks is shipped in a large crate on a pallet that weighs between 1000 lbs (454.55 kg) and 1500 lbs (681.82 kg) and measures 53" × 40" × 100.5" (1346 mm × 1016 mm × 2553 mm) (Figure 17).

Figure 17. Crate Containing One SFA10000 Rack

Each crate includes a ramp that can be used to remove the rack from the crate. A pallet jack must be used to move the crate from the delivery vehicle to a flat location that is at least 5' × 12' (1.52 m × 3.66 m) where the ramp can be installed and the rack can be rolled out of the crate. Due to the weight of the crate, it is essential that either the delivery vehicle is equipped with a lift gate or the destination has a loading dock with a lift.

Once a rack has been removed from the crate, it can be rolled to its destination. Racks are either 42U or 45U. The 42U racks are 81.75" (2077 mm) tall and the 45U racks are 87" (2210 mm) tall. Both type of racks are 28" (712 mm) wide and 45.75" (1162 mm) deep including the front and rear doors.

Note that the physical disk modules are shipped separately from the disk enclosures. The disk modules are packaged in separate boxes that are strapped to pallets. These pallets can be moved to the installation site using a pallet jack.

To Do List:

- Plan the unpacking location where the delivery vehicle will deliver the crate.
- Verify that the route to the installation site does not involve any steps.
- Verify that any elevator that is part of the route has sufficient door height and load capacity.
- Verify that the size of all the doorways along the delivery route is big enough for the pallet jack with the crate to pass through.
- Verify that there is enough space at the unpacking location to set up the ramps and remove the racks from the crates.

2.2.2 Rack Location, Air Flow, and Access

The racks that comprise a single SFA10000 storage subsystem must be installed adjacent to one another so that the disk enclosure cables can be properly routed and connected. The airflow through the racks is from front to back. The front of one rack should not be near the back of another rack.

> **NOTE :** Racks with more than 5 disk enclosures have an anti-tip plate installed at the floor level and it is important that it does not cover a cold air grill.

Access to the equipment in the racks is from their fronts and the backs. Front access is for physical disk, disk enclosure, controller chassis, and UPS maintenance. Rear access is for power supply and I/O module maintenance plus cabling.

The recommended service clearances are 48" (1220 mm) from the front of a rack and 30" (762 mm) between the back of an enclosure and a wall or 30" (762 mm) between the back of a SFA10000 system and the back of another enclosure that is exhausting hot air (Figure 18). No service clearance is required on the side of the rack.

Figure 18. SFA10000 Service Area and Dimensions

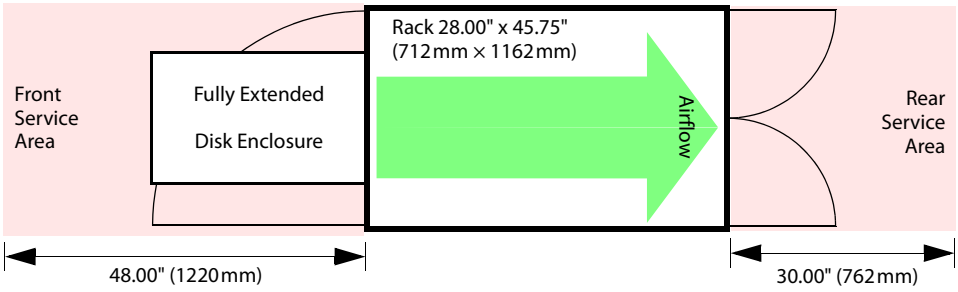

To Do List:

- Plan the location of the racks so that they have sufficient access and proper airflow. A scale drawing is recommended.

The best way to accomplish this is to create rows of racks that are side-by-side where each rack draws cool air from a 60" (1524 mm) walking isle and blows hot air into a 48" (1220 mm) maintenance isle (Figure 19).

Figure 19. Rows of Racks with Walking Isles and Maintenance Isles

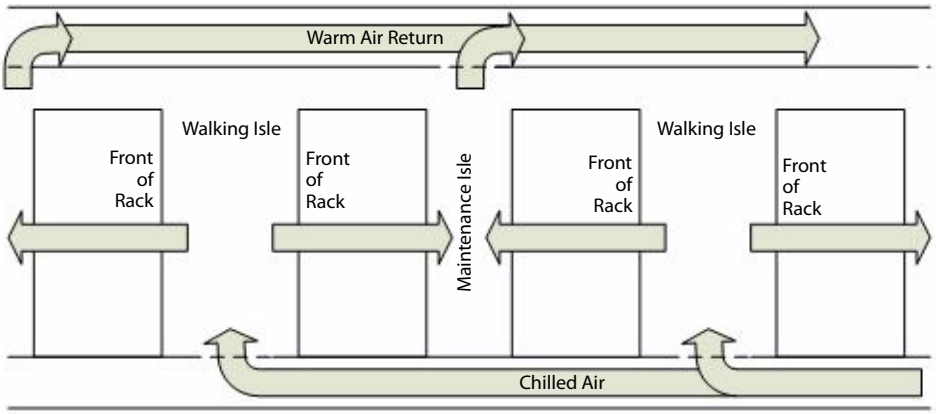

2.2.3 Floor Loading

The SFA10000 racks are heavy and it is important to verify that the weight does not exceed the floor specifications. DDN will provide an estimate of the weight of each rack on request.

A full 45U SFA10000 rack with ten SS6000 disk enclosures can weigh up to 2570 lbs (1166 kg). This weight is relatively evenly distributed across the four castors and/or leveling feet, so each will be subject to a load of 645 lbs. While the load from such a rack is 289 lb/ft^2 (2570 lb/8.9ft^2) underneath the rack, the average load is less given the space for isles. With the minimum size isles (48" in the front and 30" in the back), the load is 156 lb/ft^2 (289 lb/ft^2 × 45.75" / {45.75" + 48"/2 + 30"/2}).

The more common 42U SFA10000 racks for 5 and 10 disk enclosure subsystems weigh approximately 1600 lbs and so the point load is only 400 lbs and the floor loading is 180 lb/ft^2 (1600 lb/8.9ft^2). Accounting for the minimum isle space the load is 98 lb/ft^2.

To Do List:

- Verify with the building structural engineer that the floor structure is sufficient to hold the weight of the SFA10000 racks configured as planned in Section 2.2.2 "Rack Location, Air Flow, and Access".

2.2.4 Cooling Supply Planning

Each SFA10000 45U rack can draw up to 14.5 KW and generate 50 KBTU/hour of heat.

Each SFA10000 42U racks can draw over 8 KW and 25 KBTU/hour of heat.

Sufficient air conditioning must be provided to cool this heat load to a nominal room temperature of 25 °C. DDN will provide an estimate of power and heat for each rack on request.

To Do List:

- Verify that the rack locations will have sufficient cooling.

2.2.5 AC Power Supply Planning and Verification

Each SFA10000 rack has four to eight power cords that are designed to connect to 208VAC or 230VAC. DDN will provide a count of the power cords for each rack on request.

Each AC power cable has a NEMA L6-30P connector (Figure 20). Before the SFA10000 racks can be installed, sufficient L6-30R receptacles for these AC power cables must be installed by an electrician. Each L6-30R receptacle should be connected to a 30Amp circuit breaker.

Figure 20. L6-30R Wiring Diagram

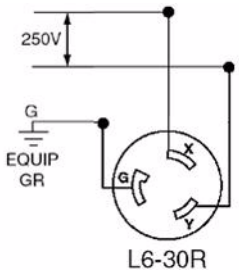

These power cords can either be connected to two separate AC power sources (for example, utility power and UPS power) or to a single AC power source. That is, half the power cords provide sufficient power to operate the system, however, all power cords must be plugged to take advantage of redundant DC power supplies.

Each SFA10000 controller has a UPS that is used to protect its cached data in the event of a power failure. These must be configured with firmware and switch settings depending on the AC voltage and how the site power is grounded. This configuration is best done at the DDN factory and so it is helpful to provide these measurements to DDN as part of the order.

To Do List:

- Measure and note down the three voltages X-Y, X-G, and Y-G as shown in Figure 20.
- Each UPS unit has two DIP switches on the back. During installation, verify that the switches are correctly set as described in Section 2.5.7.
- After power on, verify that the UPS firmware versions match the above measured voltage as described in Section 2.7.3. If these are not correct then have a DDN Field Engineer correct them.
  - If either voltage X-G or voltage Y-G is less than 10 volts then the UPS firmware version should be: Cont:01.20 Inve:01.02
  - If both voltages X-G and Y-G are greater than 10 volts then the UPS firmware version should be: Cont:01.20 Inve:83.74

2.3 Unpacking the SFA10000

The SFA10000 system components are already installed in the rack(s) if you ordered the controllers and disk enclosures with the racks.

Refer to the *DDN 42U/45U Rack Installation and Service Guide* for detailed instructions on removing the rack from the shipping crate and positioning the rack.

Refer to Appendix E "Rack Installation Procedure" if you need instructions on how to install a SFA10000 controller or UPS unit into a rack.

2.3.1 Packing List

The SFA10000 ships with the following:

- Installed in the rack(s):
  - Two controllers
  - Two UPS units
  - Four 1-meter Ethernet ICL cables
  - Two 1-meter InfiniBand ICL cables
  - Two USB cables
  - Forty SAS cables
  - Five, ten, or twenty disk enclosures
- Disk modules
- Two serial cables
- Two Ethernet cables
- SFA10000 Quick Start Guide
- SFA10000 User Guide
- 42U/45U Rack Installation and Service Guide
- SFA OS CLUI Command Reference
- Release Notes

2.4 Installing the Disk Modules

The disk modules are shipped separately from the disk enclosures.

The SFA10000E disk modules have been pre-labeled with a slot number. Use the charts in Appendix C to determine proper drive placement according to the slot number on the label.

For SFA10000 configurations, create a more balanced load by evenly distribute the disk modules among the disk enclosures. If a mixture of disk technologies, such as SAS and SATA, will be populated into the enclosures, it is best to populate the SAS disks into the front slots and install the SATA disks in the rear slots.

Wear an ESD wrist strap or otherwise ground yourself when handling the disk modules and components. Electrostatic discharge can damage the circuit boards.

Follow these steps to install a disk module:

NOTE : Also refer to Appendix B if you need detailed instructions for disk module installation.

1. For system using SS6000 or SS7000 disk enclosures:
   - Slide the enclosure out from the rack by squeezing the tabs on both rack slides. Keep pulling until the enclosure locks and you hear a clicking sound.
   - Open the enclosure covers.
2. Release the handle on the disk module.
3. Insert the module into a disk bay. Cam the disk module home. The camming foot on the base of the module will engage into the slot in the enclosure.
4. When the module is fully inserted, close the handle. You should hear a click as the latch engages and holds the handle closed.

NOTE : On the SS2460 enclosures, a disk module or a dummy disk module must be installed in every slot of the enclosure in order to maintain proper airflow and cooling.

5. For system using SS6000 or SS7000 disk enclosures:
   - After you have installed all the disk modules in this enclosure, close the enclosure covers and engage both cover latches.
   - Squeeze the tabs on both rack slides and push the enclosure back into the rack.

2.5 Cable Connections

2.5.1 Couplet ICL Cabling

There are two sets of Inter-Controller Link (ICL) connections between the two controllers.

1. Verify that the two InfiniBand cables are attached to the ICL ports on the two controllers as shown in Figure 21 for SFA10000 and Figure 22 for SFA10000E.

2. Verify that the four short Ethernet cables are attached to the Ethernet ports on the two controllers as shown in Figure 21 for SFA10000 and Figure 22 for SFA10000E.

Figure 21. Inter-Controller Link Connections on SFA10000 Controllers

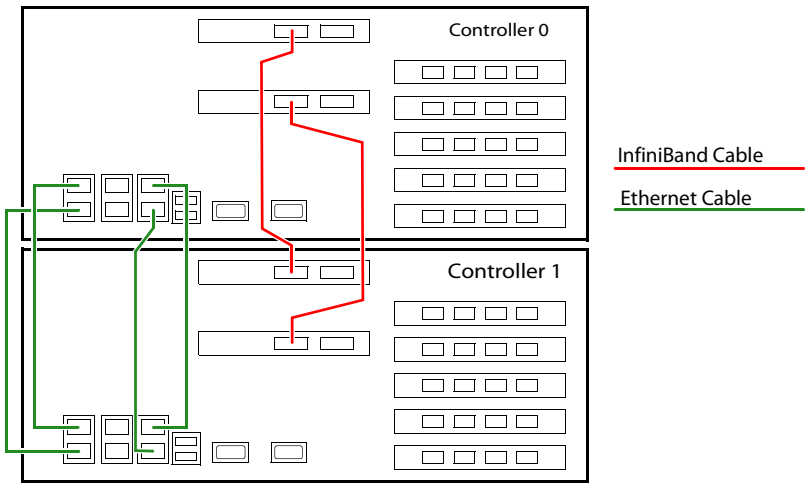

Figure 22. Inter-Controller Link Connections on SFA10000E Controllers

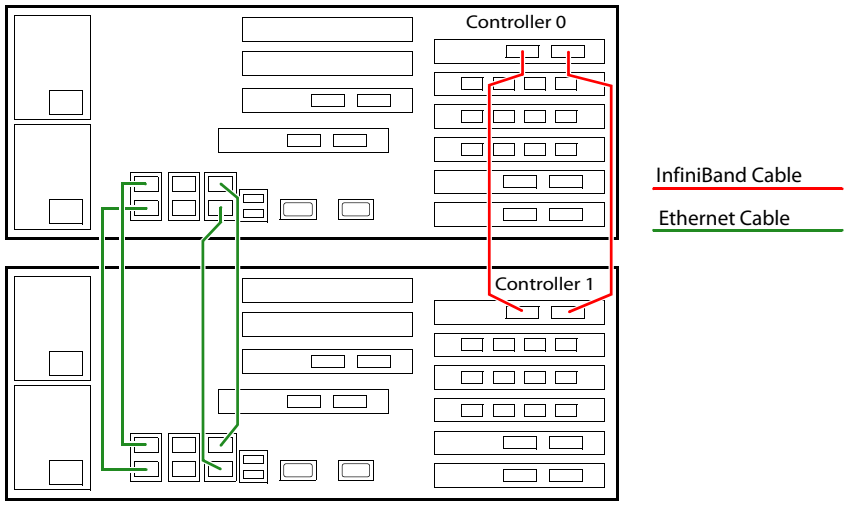

2.5.2 Disk Enclosure Cabling

The SFA10000 systems are shipped with the cables attached between the disk enclosures and controller's I/O channels.

NOTE : Incorrect wiring can prevent the system from operating correctly or from operating at all.

Verify that the disk enclosures are correctly connected to the two controllers using the figures listed in the tables below. The cables are labeled by ports to which they will be connected. In the illustrations, the port numbers are colored to aid in locating matching cables.

SFA10000 configurations using SS6000 enclosures:

Figure 23	Five enclosures
Figure 24	Ten enclosures
Figure 25	Twenty enclosures

SFA10000 configurations using SS7000 enclosures that contain either SAS disks, NearLine SAS disks, or SATA physical disks with a SAS-to-SATA bridge interposers:

Figure 23	Five enclosures
Figure 24	Ten enclosures
Figure 25	Twenty enclosures

SFA10000 configuration using SS2460 enclosures:

Figure 26	Five enclosures containing SAS disks only

SFA10000E configurations:

Figure 27	Five SS6000 disk enclosures
Figure 28	Ten SS6000 disk enclosures

Figure 23. Connecting SFA10000 Controllers to 5 x SS6000 Enclosures or 5 x SS7000 Enclosures

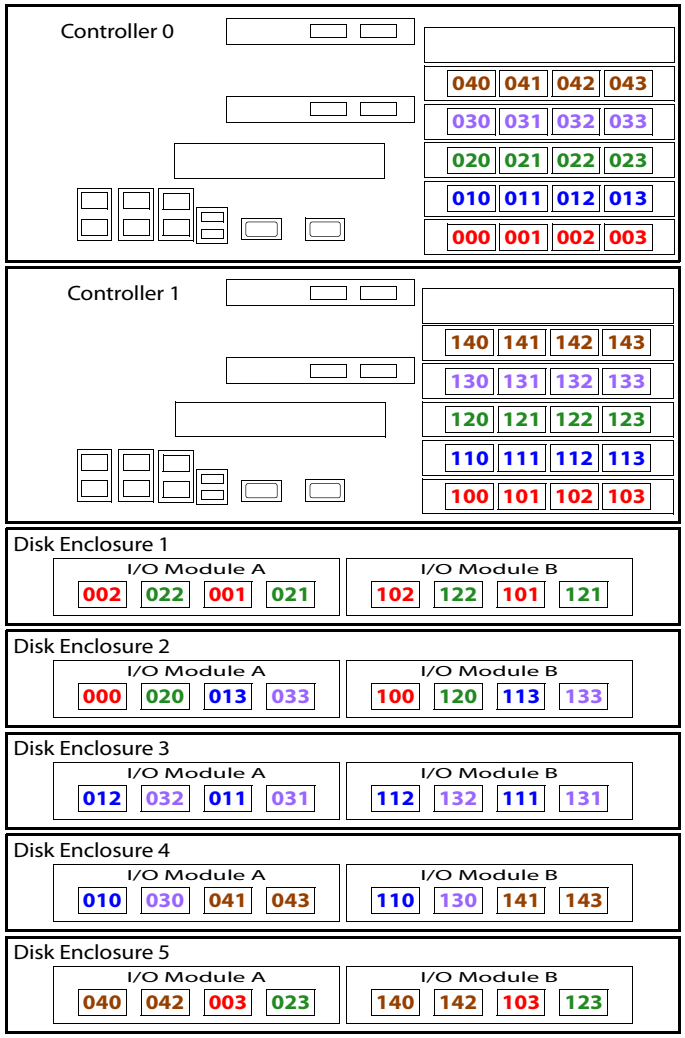

Figure 24. Connecting SFA10000 Controllers to 10 x SS6000 Enclosures or 10 x SS7000 Enclosures

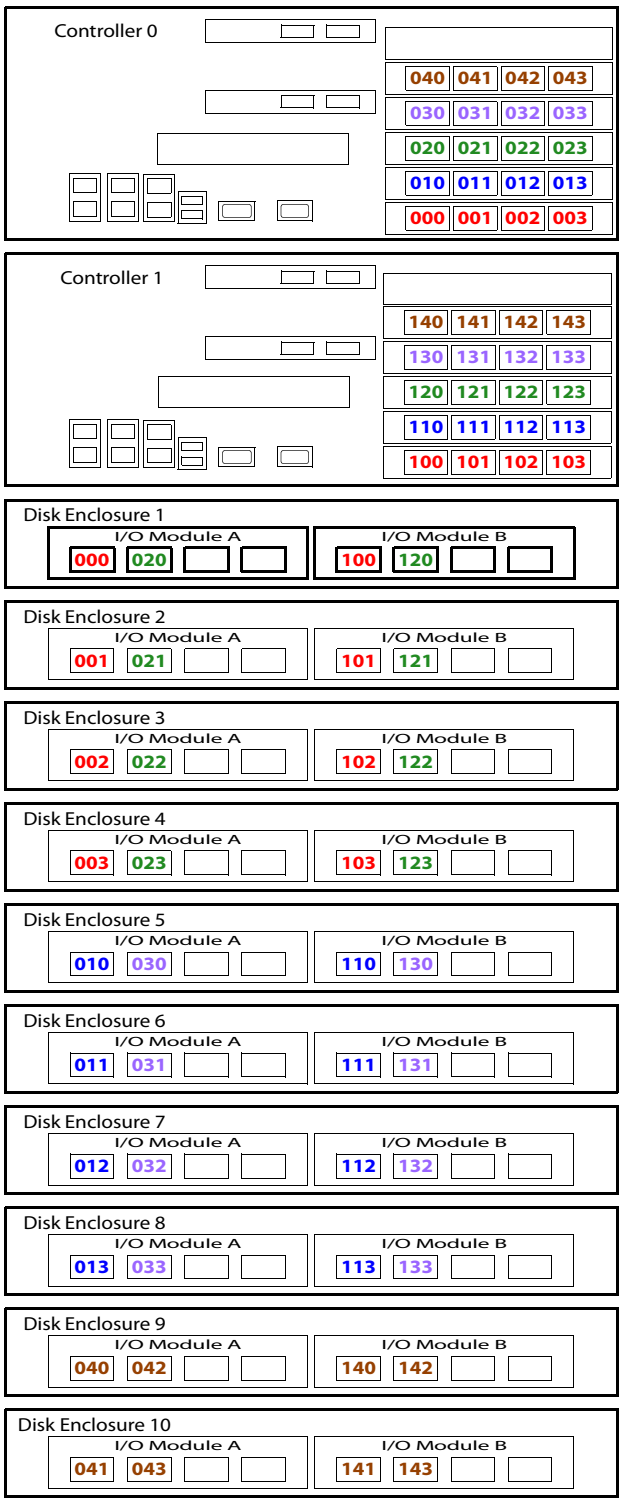

Figure 25. Connecting 20 x 60-Bay Enclosures to SFA10000 Controllers

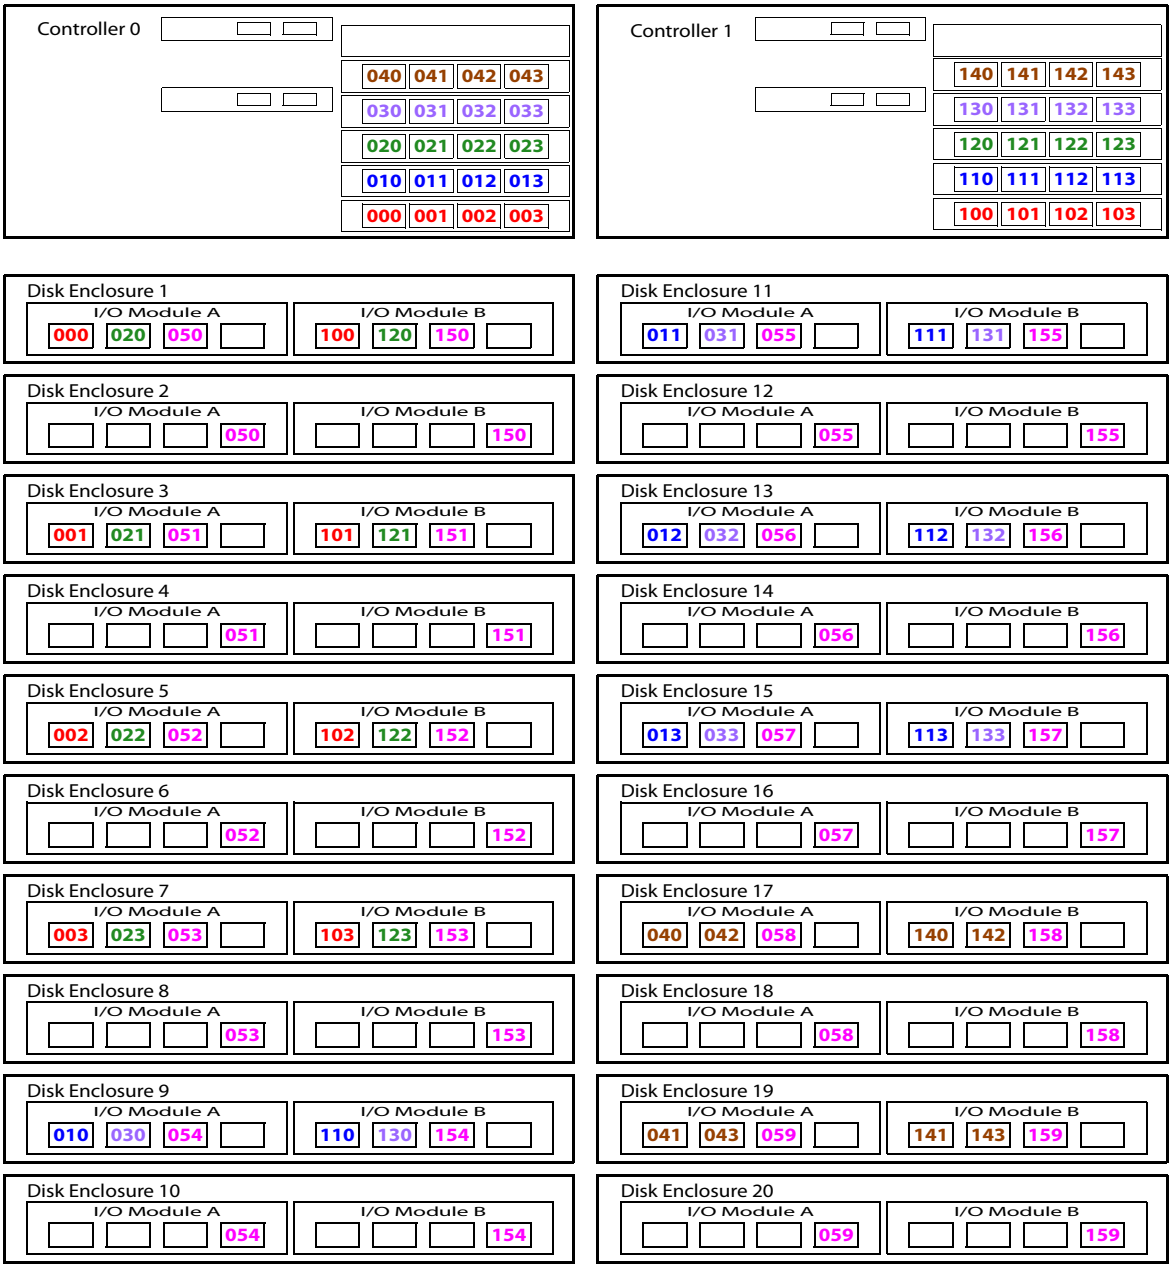

Figure 26. Connecting 5 x SS2460 Enclosures to SFA10000 Controllers

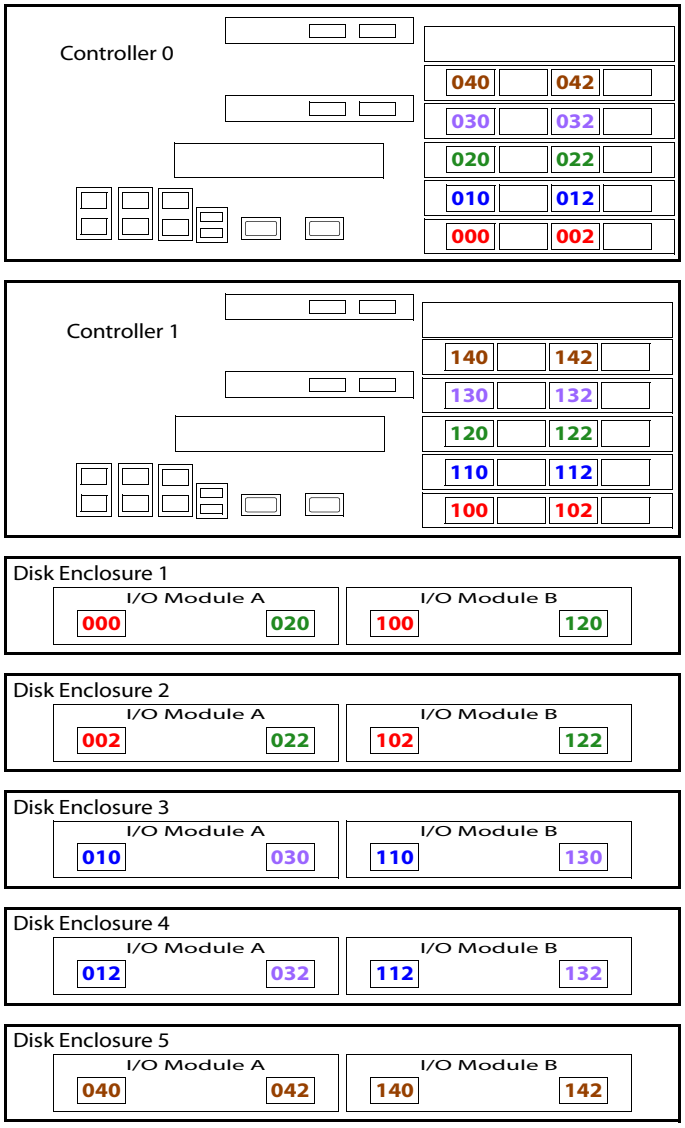

Figure 27. Connecting 5 x SS6000 Enclosures to SFA10000E Controllers

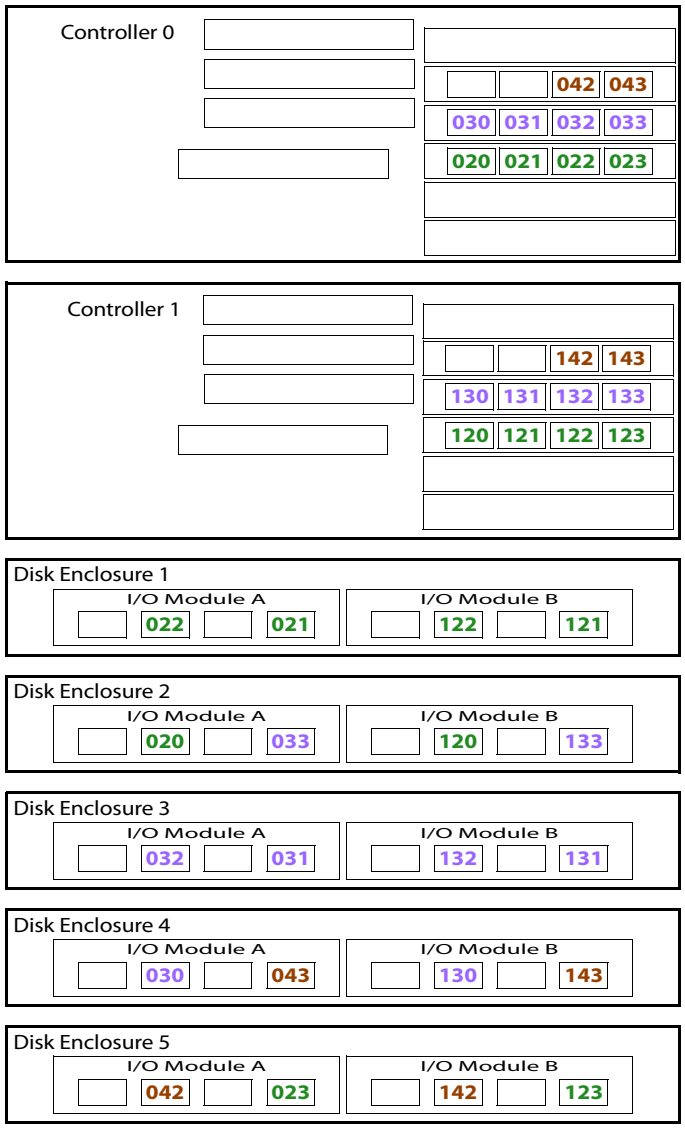

Figure 28. Connecting 10 x SS6000 Enclosures to SFA10000E Controllers

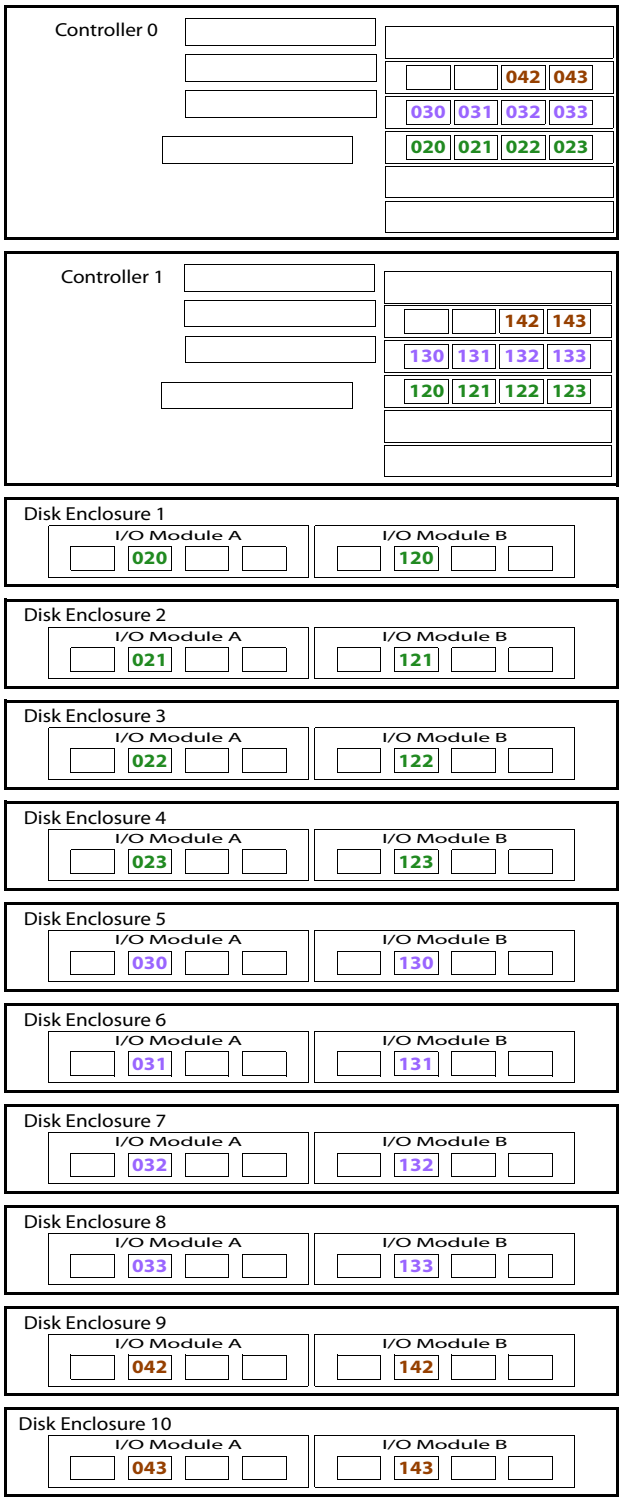

2.5.3 Host Connections

2.5.3.1 SFA10000

Connect the controller host ports shown in Figure 29 either directly to your hosts or to a switch that connects your hosts. Depending on your SFA10000 model, these connections may be Fibre Channel or InfiniBand.

NOTE : Do NOT use the unused Inter-Controller Link connectors to connect InfiniBand hosts or switches.

Figure 29. Host Ports on SFA10000 Controllers

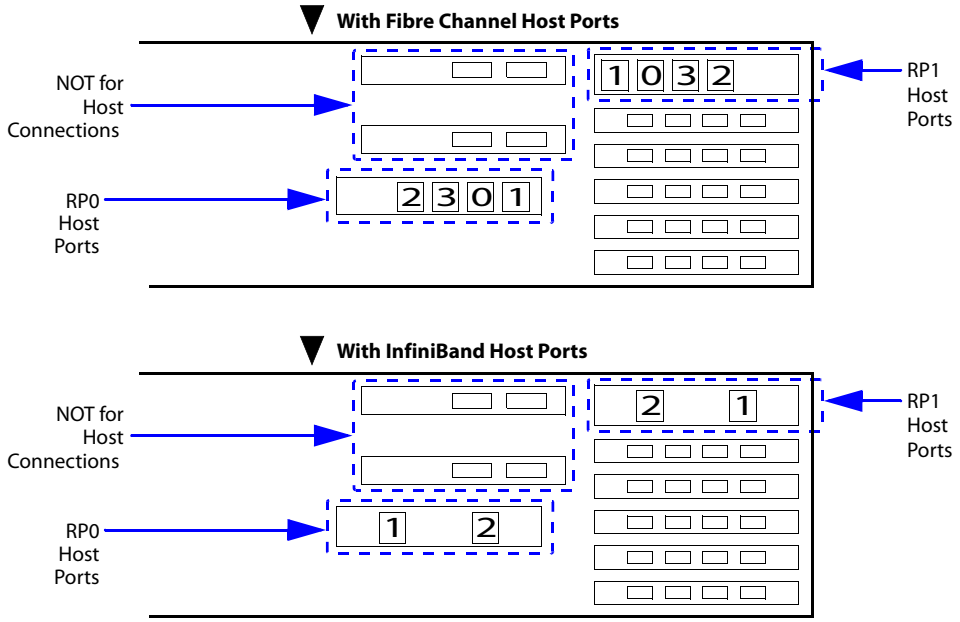

2.5.3.2 SFA10000E

Connect the controllers' host ports either directly to your file system clients or to a switch that connects your clients together. The SFA10000E supports InfiniBand (IB), 10G Ethernet, or a combination of the two. The type of cable used depends upon your client network and application needs. The required cables are shipped with your SFA10000E. If not, please contact DDN Technical Support.

The host ports are numbered as shown in Figure 30. For each host card on the controller, the cables can be connected in one of the following configurations:

	Port 1	Port 2
1)	IB	IB
2)	10Ge	10Ge
3)	IB	10Ge

NOTE : A combination of IB and 10Ge cables can only be used in one configuration—#3 above.

Figure 30. Host Ports on SFA10000E Controllers

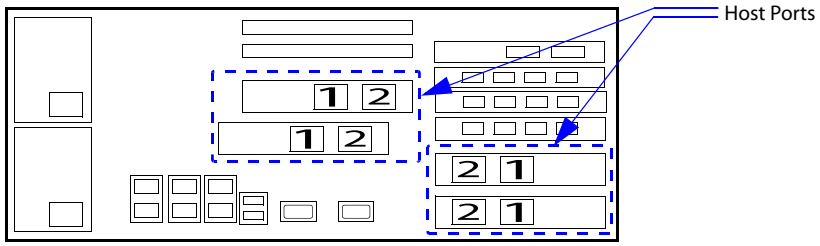

2.5.4 Management Network Connection

You may remotely monitor the system over your Ethernet network. Connect the controllers to your network using the Ethernet port as shown in Figure 31. The Ethernet management port supports 10, 100, and 1000BASE-T rates.

Figure 31. Ethernet Connections to Your Network

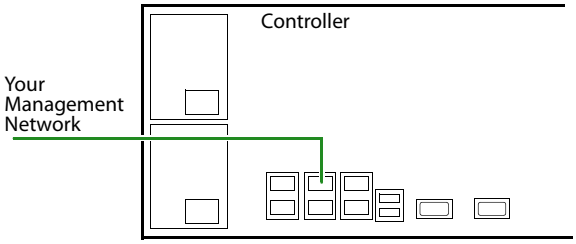

2.5.5 RS-232 Console Connection

Connect a null modem cable between a PC and the RS-232 connector on the back of the controller (Figure 32). The RS-232 console is used to configure the management network ports during initial configuration. It can also be used to log the console output and upgrade the BIOS/BMC/CONFIG firmware.

Figure 32. Connect RS-232 Console to Controller

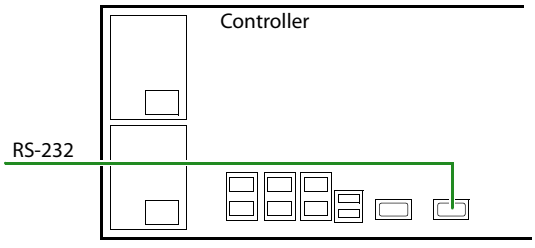

2.5.6 UPS Connection

1. For each controller, verify that a USB cable is attached to the port on the back of the controller and its UPS as shown in Figure 33.
2. For each controller, verify that the two power cables are attached to the controller's power supplies and its UPS as shown in Figure 33.

Figure 33. Connect UPS to Controller

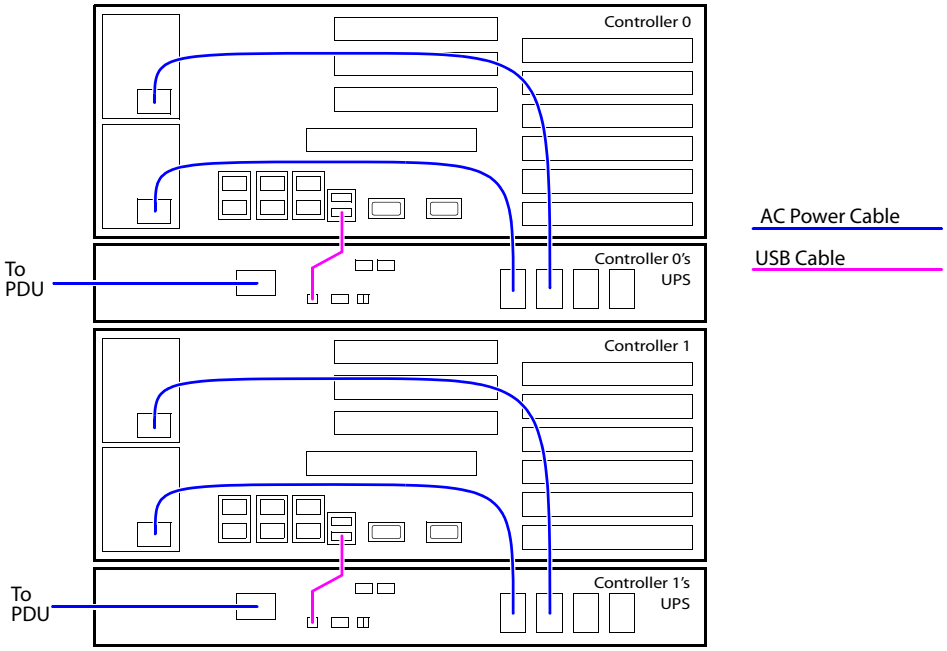

2.5.6.1 UPS Battery

The UPS units are shipped with the batteries disconnected. Follow these steps to connect the batteries:

1. Remove the front bezel of the unit.
2. Fasten the two connectors together (Figure 34). ***Note that it is normal if you see a spark and hear a pop sound as you connect the battery.***
3. Replace the front bezel.
4. Repeat the above steps on the other UPS.

Figure 34. Connect UPS Battery

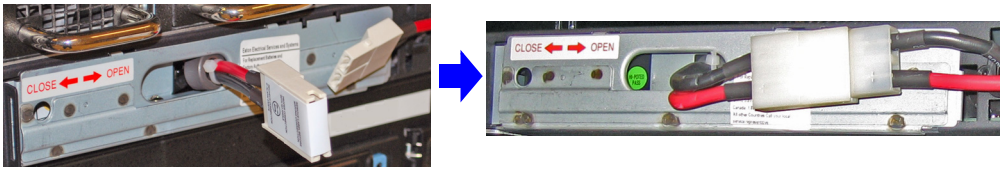

2.5.7 Power Connections

1. Verify your AC power source by measuring the AC voltage.
2. Refer to the illustrations in Figure 35 and verify that the DIP switches on the back of both UPS units are correctly set.

NOTE : Changing the voltage of the UPS requires more steps than just changing the DIP switches. If the settings are incorrect, have a DDN field engineer correct them.

Figure 35. DIP Switch Settings on UPS Rear Panel

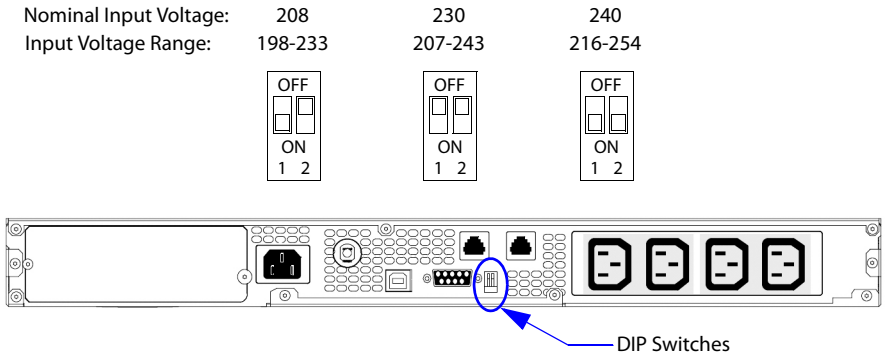

3. Verify that the power switches on all the power distribution units (PDUs) are set to OFF.
4. Connect the PDUs to your AC power source. For maximum redundancy, connect the PDUs to different AC circuits.

2.6 Powering On the System

1. Switch on all the PDUs.
2. Switch on the power supply modules on all the disk enclosures.
3. On both UPS units, press and hold the Power button for 2 seconds on the front panel to turn on the units (Figure 36).

Figure 36. UPS Front Panel

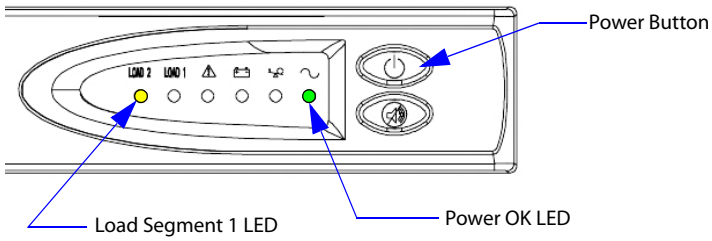

4. Verify that the Power OK LED on both UPS units turns green and the Load Segment 1 LED turns yellow (Figure 36), indicating a successful power application.
5. Verify that the Power indicator on both controllers turns green (Figure 37). If not, press the Power button once.
6. Verify that the Fault LED on both controllers is off (Figure 37).

Figure 37. Controller Front Panel LEDs

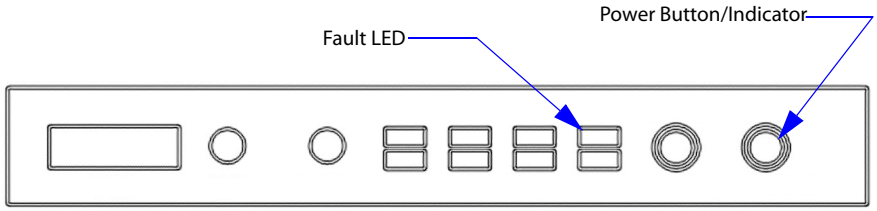

You may now configure the system as described in Section 2.7 "Configuring the SFA10000".

2.7 Configuring the SFA10000

This section provides information on configuring your SFA10000.

NOTE : The configuration examples provided here represent only a general guideline. These examples should not be used directly to configure your particular SFA10000.

The CLUI (Command Line User Interface) commands used in these examples are fully documented in Chapter 3—however, exact commands may change depending on your firmware version. To access the most up-to-date commands, use the CLUI's online HELP feature.

The CLUI commands are independent of case. Most of the keywords can be abbreviated and most of the punctuations are optional. For example, "SHOW CON LO ALL" is adequate for "SHOW CONTROLLER LOCAL ALL_ATTRIBUTES".

2.7.1 Planning Your Setup and Configuration

Before proceeding to configure the storage settings for the SFA10000, it is necessary to understand the basic organization of the system.

The SFA10000 uses storage pools and virtual disks to configure disk storage for use by the hosts. A storage pool (sometimes referred to as a RAID group) is a collection of 2 to 10 physical disks, ideally of the same capacity and type. In RAID 5, the capacity of one disk is reserved for parity, allowing data recovery in the event of a single disk loss in the storage pool. In RAID 6, the capacity of two disks is reserved for parity, allowing data recovery if either one or two disks are down in a storage pool. In RAID 1, the capacity of one disk is used for data duplication.

A single RAID 1 storage pool can be configured with 2 disks. A single RAID 5 storage pool can be configured using 5 or 9 disks. A single RAID 6 storage pool can be configured using 6 or 10 disks.

NOTE : It is **NOT** recommended to configure storage pools with physical disks of different capacity as the SFA10000 will use the lower capacity for all disks and space will be wasted.

A virtual disk is the storage unit presented to any attached host. Virtual disks allocate space in 8 GiB increments. For example, 16 GiB of storage space will be allocated when creating a virtual disk of 10 GiB.

The disks in the system are categorized into one of the following pools:

- Unassigned Pool - By default, all the disks or replacement disks are initially allocated to this pool. When a disk has failed or been released (from a spare pool, for example), it is also moved to this pool.
- Storage Pool (sometimes referred to as a RAID group) - A collection of 2 to 10 disks, ideally of the same capacity and type.
- Spare Pool - This type of pool contains disk(s) that can be used as spare disks in one or more storage pools.

2.7.2 Serial Interface Configuration

1. On the PC that is connected to the controller, load a serial console program (such as HyperTerminal, minicom, and screen) and use the following settings for the serial connection:
   - Bits per second: 115,200
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: None
2. Once connected, press the <Enter> key to bring up the login prompt. Enter the user name **`user`** and password **`user`**.

2.7.3 Validate the Hardware

1. Enter: **`SHOW ENCLOSURE * ALL`**

Figure 38. Show Enclosure Status

```
************************
*     Enclosure(s)     *
************************
Index:                    0
Enclosure:                0x50000000
Type:                     CONTROLLER
Logical ID:               0x0001ff0801c90000
Vendor ID:                DDN
Product ID:               SFA10000
Revision:                 0000
Zones:                    0
Fault Indicator:          OFF
Predicted Failure Ind:    OFF
Locate Indicator:         OFF
Baseboard part number:    B48-DTIAAD6G00D100
Baseboard serial number:33759-1017-00200
Baseboard date:           20091005
Chassis part number:      X3B70W50A81XXXXXXP
Chassis serial number:    516-100427003000
BIOS version:             GEMDV233
BMC version:              DDGEN050
BMC config:               JANUSC21
Firmware version:         1.4

Index:                    1
Enclosure:                0x50000001
Type:                     DISK
Logical ID:               0x50001ff210030000
Vendor ID:                DDN
Product ID:               SS7000
Revision:                 0402
Zones:                    2
Fault Indicator:          OFF
Predicted Failure Ind:    OFF
Locate Indicator:         OFF
Baseboard part numbmer:   USE13000098RB05D
Firmware version:         04.02.01
FPGA version:             0008

Index:                    2
Enclosure:                0x50000002
Type:                     DISK
Logical ID:               0x50001ff210006000
Vendor ID:                DDN
Product ID:               SS7000
Revision:                 0402
Zones:                    2
Fault Indicator:          OFF
Predicted Failure Ind:    OFF
Locate Indicator:         OFF
Baseboard part number:    USE13000098RB05D
Firmware version:         04.02.01
FPGA version:             0008
 :
 :
```

Controller Firmware Version

Zoning Configuration

Disk Enclosure Firmware Version

Zoning Configuration

Disk Enclosure Firmware Version

- Verify that both controllers and all the disk enclosures can be seen (Figure 38) and contain consistent firmware versions.

NOTE : If an enclosure is not reporting the firmware version or reporting an incorrect version of firmware, you must contact DDN Technical Support to have this error corrected before proceeding with the configuration.

- Verify that the disk enclosures are correctly zoned (Figure 38) using the table below. If the setting is incorrect on a SS6000 enclosure, refer to Appendix D for instructions on changing the configuration. If the setting is incorrect on a SS7000 or SS2460 enclosure, then have a DDN field engineer correct it.

Zones: 1	• Systems using ten 60-bay disk enclosures • Systems using twenty 60-bay disk enclosures • Systems using five SS2460 disk enclosures
Zones: 2	• Systems using five 60-bay disk enclosures

2. Enter: **SHOW PHYSICAL_DISK ***
Verify that all the disks can be seen and are healthy (Figure 39). If the disks are not visible, type **SHUTDOWN SUBSYSTEM RESTART** to reboot the controllers.

Figure 39. Show Physical Disk Status

```
RAID[0]$ show pd *
                                  |Health|
Encl|Slot|Vendor| ::::::::::::::::State|Idx|State|::::::
-------------------------------------------------------------
   1    1 HITACHI :::::::::::::::: GOOD    6 NORM  :::::
   1    2 HITACHI :::::::::::::::: GOOD    1 NORM  :::::
   1    3 HITACHI :::::::::::::::: GOOD    5 NORM  :::::
:::::::::::::::::::::::::::::::::::::::::::::::::::::::::::

Total Physical Disks:   300
Total Assigned Disks:   0
Total Unassigned Disks: 300
  Total SAS Disks:      300
  Total Member State Normal:  300
```

3. Enter: **SHOW FAN ***
Verify that no fan faults are found.

4. Enter: **SHOW POWER ***
Verify that the “AC Fail” and “DC Fail” states are false.

5. Enter: **SHOW TEMPERATURE ***
Verify that all the temperature sensors do not report over temperature conditions.

6. Enter: **SHOW UPS * ALL**
Verify that the UPS batteries are healthy (Figure 40). Also take note of the charge level as the battery may take several hours to fully charge.

   Verify that the UPS firmware versions match the measured voltages noted down in Section 2.2.5 "AC Power Supply Planning and Verification". If these are incorrect, then have a DDN field engineer correct them.

   - If either voltage X-G or voltage Y-G is less than 10 volts then the UPS firmware version should be: Cont:01.20 Inve:01.02
   - If both voltages X-G and Y-G are greater than 10 volts then the UPS firmware version should be: Cont:01.20 Inve:83.74

Figure 40. Show UPS Attributes

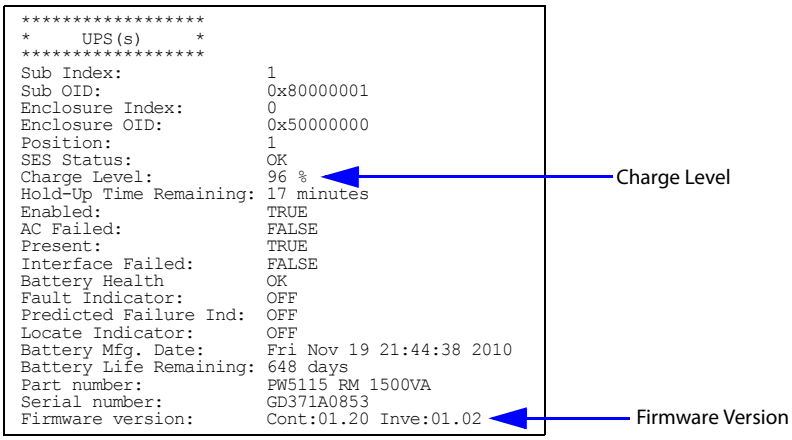

2.7.4 Clear System Configuration (SFA10000 only)

Warning

Use of the CLEAR SUBSYSTEM CONFIGURATION command will destroy any existing data.

NOTE : For SFA10000E, the storage has been preconfigured, so the existing configuration should **NOT** be cleared. Skip this step and proceed to Section 2.7.5 "Set the System Name".

To ensure that you are starting from a fresh/clean install, enter:

CLEAR SUBSYSTEM CONFIGURATION

Then enter **YES** twice to confirm deletion of current configuration (Figure 41).

Figure 41. Clear Subsystem Configuration

```
RAID[0]$ clear subsystem configuration
Are you sure you want to delete this configuration [No]? yes
Are you sure you want to delete this configuration?
All data will be lost [No]? yes
RAID SUBSYSTEM CONFIGURATION cleared STATUS='Success' (0x0)
```

2.7.5 Set the System Name

You may set the system name using the command (Figure 42):

SET SUBSYSTEM NAME=<name>

To verify the new setting, enter the command:

SHOW SUBSYSTEM ALL_ATTRIBUTES

Figure 42. Set Subsystem Name

```
RAID[0]$ set subsystem name=SFA10000
SUBSYSTEM attributes set STATUS='Success' (0x0)

RAID[0]$ show subsystem all
RP Subsystem Name:      SFA10000
UID:                    60000000000000000000000030000000
 .
 .
 .
```

2.7.6 Set System Time & Date (NTP)

NTP (Network Time Protocol) mode is available on the SFA10000. You can enter up to four NTP addresses as the time servers.

To enable the NTP mode, enter the command:

SET SUBSYSTEM NTP=[<ip address list, up to 4>|NONE]

Specifying a list of NTP addresses will start the NTP mode on each controller, using that list of NTP addresses as the time servers. Specifying **NONE** will turn off NTP mode.

To display the current settings, enter the command (Figure 43):

SHOW SUBSYSTEM ALL_ATTRIBUTES

Figure 43. Show NTP Settings

```
********************
*    Subsystem     *
********************

RP Subsystem Name:      SFA10000
UID:                    60001ff0800a30000000000030000000
Subsystem Time:         Fri Apr 22 6:26:39 2011
Locate Dwell Time:      120 seconds
Fast Timeout:           ON
Pool Verify Priority:   10%
NTP Mode:               ON
                        (10.32.16.24)
                        (10.32.16.25)
                        (10.32.16.26)
                        (10.32.16.27)
Drive Error Tolerance:  HIGH
Single Controller WB:   ENABLED
Fri Apr 22 6:26:39 2011
```

If you are not using NTP, you can set the time of both controllers using the **SET SUBSYSTEM DATE_AND_TIME** command. However, once the time is set, the time on the two controllers is free to drift independent of the other controller. To change the date and time to March 1, 2011 2:15 pm, for example, type: **SET SUBSYSTEM DATE_AND_TIME=2011:3:1:14:15**

NOTE : If NTP mode is enabled, the **SET SUBSYSTEM DATE_AND_TIME** command will fail with a status of "Setting date/time while in NTP mode".

2.7.7 Configure Network Interface Settings

The SFA10000 can be configured and administered either via serial connection (using the supplied serial cable) or via Ethernet connection using SSH. However, in order to use the SSH connection, it is first necessary to configure the network settings on each controller. This can only be done using the serial interface as described below.

1. Enter the command:
   UI SET NETWORK_INTERFACE=LOCAL 0 IP_ADDRESS=<ip_address> IP_MASK=<netmask> IP_GATEWAY=<gateway>
   where <ip_address> is an address appropriate to the local network. The netmask and gateway values can also be entered(Figure 44).

Figure 44. Set Network Interface Example

```
RAID[0]$ ui set network_interface=local 0 ip_address=192.168.0.10
ip_mask=255.255.255.0 ip_gateway=192.168.0.1
NETWORK_INTERFACE 0 set with STATUS='Success' (0x0)
```

2. To verify the new settings, enter the command (Figure 45):
   UI SHOW NETWORK_INTERFACE=LOCAL

Figure 45. Show Network Interface Settings

```
RAID[0]$ ui show network_interface local

Network device id 0
    address 192.168.0.10
    netmask 255.255.255.0
    gateway 192.168.0.1
```

3. Switch the serial cable to the RS-232 port on second controller and repeat the above steps to enter a different IP address for second controller.

At this point, you should be able to use the SFA OS CLUI and GUI via the network.

If you want to set up email and SNMP event notification, refer to Section 3.13.5 "Email and SNMP Notification Setup" on page 102 for information.

2.7.8 Access Virtual Machines (SFA10000E Only)

The SFA10000E is shipped with Virtual Machines (VM) preconfigured and installed for serving your ExaScaler or GridScaler file system.

NOTE : By default, each ExaScaler system is shipped with three VMs enabled and a GridScaler system is shipped with four VMs enabled.

The Virtual Machines can be accessed in one of two ways:

- **VNC**
  You can use a VNC viewer to connect to a VM console. Use the IP address followed by a colon, then followed by a single digit in the range 0..3. The last digit represents the VM you wish to connect to. For example, if the IP address of the controller is 192.168.0.1 and you want to connect to the first VM, use a VNC connection address of "**192.168.0.1:0**".

- **SSH**
  You can use a standard terminal program with SSH support to connect to a VM console. This is done by specifying a specific port to the SSH connection address. The valid ports are 5555..5558—each representing the VM you wish to connect to. For example, if the IP address of the controller is 192.168.0.1 and you want to connect to the first VM using a standard Linux shell, enter "**root@192.168.0.1:0**".

NOTE : The default password for the root user is **password**. This should be changed immediately on each VM.

The SFA10000E initial configuration is complete at this point.
The procedures given in Sections 2.7.9 through 2.7.13 can be skipped.
With root access to each Virtual Machine, you are now able to configure your ExaScaler or GridScaler file system.

2.7.9 Create Storage Pools

A storage pool on a SFA10000 has the following attributes:

- RAID Level (RAID)
  Storage pools can be configured to use either a RAID 1, RAID 5 or RAID 6 parity scheme. In RAID 1, the capacity of one disk is used for data duplication. In RAID 5, the capacity of one disk is reserved for parity, allowing data recovery in the event of a single disk loss in the storage pool. In RAID 6, the capacity of two disks is reserved for parity, allowing data recovery if either one or two disks are down in a storage pool. ***For maximum data protection, DataDirect Networks recommends the use of RAID 6***.

- Chunk Size (CHUNK)
  The chunk size (in KiB blocks) defines the amount of data written to a single disk before proceeding to the next disk in the storage pool.

NOTE : RAID 1 is a two member RAID set where the data is mirrored on each disk. There is no parity, hence, the chunk size is fixed.

- Disk Count (NUMBER)
  A RAID 1 storage pool may consists of 2 physical disks. A RAID 5 storage pool may consist of 5 or 9 physical disks. A RAID 6 storage pool may consist of 6 or 10 physical disks. For maximum performance, select disks with the same characteristics (such as SAS/SATA, capacity, and RPM).

To create a storage pool, use the **CREATE POOL** command:

1. At the CLUI prompt, type:

```
CREATE POOL RAID_LEVEL=[raid1|raid5|raid6]
CHUNK_SIZE=[32|64|128|256] {ASSIGN_POLICY=[SAS|SATA]
NUMBER=[2|5|6|9|10] or PHYSICAL_DISK=<list of disks>}
INIT_ALLOW=<initialization policy>
SATASSURE=[NONE|DATA_INTEGRITY_FIELD|PARITY]
```

   where <list of disks>=index name of disks; <initialization policy>=Allow_IO, NO_IO, Priority (1-99, where 99 puts all system resources on the initialization).

   You may either explicitly select the disks for the pool by using "**PHYSICAL_DISK=**" or specify the number of disks in the pool by using "**NUMBER=**" in which case the next available disks will be selected (Figure 46). You may disable the SATAssure feature or enable it to use either the data integrity field option or the parity check on read option.

   Figure 46. Create Storage Pool Example (1)

```
RAID[0]$ create pool raid_level=raid5 chunk_size=64kb physical_disk=0x6c, 0x6d, 0x6e, 0x6f, 0x70
POOL 0 OID=0x19b60000 create STATUS='Success' (0x0)

RAID[0]$ create pool raid_level=raid6 chunk_size=128kb number=6
POOL 1 OID=0x19b80001 create STATUS='Success' (0x0)
```

   If you specify the **ASSIGN_POLICY** and **NUMBER** parameters together (Figure 47), only the selected type of disks will be used for the storage pool and you will not need to enter the individual disk name.

Figure 47. Create Storage Pool Example (2)

```
RAID[0]$ create pool raid_level=raid6 chunk_size=128kb number=6 assign_policy=sas
POOL 2 OID=0x19b80002 create STATUS='Success' (0x0)
```

2. Use the **CREATE POOL** command to add more storage pools as needed.
If you need to delete a storage pool, use the **DELETE POOL <pool-id>** command.

NOTE : Storage pool initialization is a background process and multiple storage pools can be initialized simultaneously. However, virtual disks on a storage pool are NOT accessible until its initialization is complete.

You may assign a name to a storage pool using the command:
SET POOL=<pool-id> NAME=<pool name>

To view the list of configured storage pools, enter the **SHOW POOL *** command (Figure 48).

Figure 48. Show Pool Information Screen

```
RAID[0]$ show pool *

                                          |Total |Free  |Max   |                  |Disk|  Global   |Spare |
Idx|Name       |State      |Chunk|Raid| Faults |cap GB|cap GB|VD GB | Settings  | Jobs |T/O |spare pool|Policy|
-----------------------------------------------------------------------------------------------------------
  0 pool-0      NORMAL       128   5              4104      0      0   W   R F I          10  UNASSIGNED  AUTO
  1 pool-1      NORMAL       128   5              4104      0      0   W   R F I          10  UNASSIGNED  AUTO
  2 pool-2      NORMAL       128   5              4104      0      0   W   R F I          10  UNASSIGNED  AUTO
  3 pool-3      NORMAL       128   5              6192      0      0   W   R F I I        10  UNASSIGNED  AUTO

Total Storage Pools: 4
```

To display the detailed information of a storage pool, use the **SHOW POOL <id> ALL_ATTRIBUTES** command (Figure 49).

Figure 49. Show Pool All_Attributes Example Screen

```
*****************
*    Pool(s)    *
*****************
  Index:                0
  OID:                  0x19e50000
  Type:                 STORAGE
  Name:                 pool-0
  Chunk Size:           128KB (0x100 blocks)
  Block Size:           512
  RAID Type:            RAID5
  Free Raid5 Capacity:  0 GB
  Max VD Size:          0 GB
  Total Capacity:       600 GB
  UUID:                 60000000000000000000180d0000
  Global Spare Pool:    UNASSIGNED
  DiskTimeout(FRT):     10 minutes
  Init Policy:          ALLOW_IO
  Init Priority:        50%
  Verify Priority:      70%
  Full Rebuild Priority:        80%
  Fractional Rebuild Priority:  90%
  Sparing Policy:       AUTOMATIC
  Verify Policy:        DISABLED
  Assign Policy:
    Device Type:        SATA
    Rotation Speed:     NA
    Raw Capacity:       NA
  SATAssure:            None
  Cache Settings:
    ReACT:              FALSE
    IO Routing:         TRUE
    Mirroring:          TRUE
    Read Ahead:         TRUE
    Write Back:         FALSE
  Initializing:         TRUE
  Rebuilding:           FALSE
  Paused:               FALSE
  AutoWriteLock:        FALSE
  Critical:             FALSE
  Forced Write-Thru     FALSE
  Current Home:         0x38000000 0x00000000 (LOCAL)
  Future Home:          0xffffffff  0x00000000
  Preferred Home:       0x38000000 0x00000000 (LOCAL)
  Job OID:              0x28000003
  Total Phy Disks       5
  State:                NORMAL
  Member Size:          120 GB
    pID     State     UUID
    0x0001  NORM      0x5000cca215c56e47
    0x0002  NORM      0x5000cca215c56456
    0x0003  NORM      0x5000cca215c54c71
    0x0004  NORM      0x5000cca215c5675c
    0x0005  NORM      0x5000cca215c56e55
```

2.7.10 Create Virtual Disks

A virtual disk can be created to use all or a part of a storage pool.

To create a virtual disk, use the **CREATE VIRTUAL_DISK** command:

1. At the CLUI prompt, type:
   CREATE VIRTUAL_DISK POOL=<pool-id> CAPACITY=<capacity>

   where <pool-id> is the storage pool to be used by this virtual disk; <capacity> is the capacity of the virtual disk in GiB or type “max” to use all available capacity (Figure 50).

   A message is displayed to indicate whether the virtual disk creation was successful.

Figure 50. CREATE VIRTUAL DISKS Example Screen

```
RAID[0]$ create virtual_disk pool=0 capacity=16
VIRTUAL_DISK 0 OID=0x89ba000 creation STATUS='Success' (0x0)

RAID[0]$ create virtual_disk capacity=32 pool=0
VIRTUAL_DISK 1 OID=0x89bb001 creation STATUS='Success' (0x0)
```

2. Use the **CREATE VIRTUAL_DISK** command to add more virtual disks as needed.

   If you need to delete a virtual disk, use the **DELETE VIRTUAL_DISK <virtual disk-id>** command.

You may assign a name to a virtual disk :
SET VIRTUAL_DISK=<virtual disk-id> NAME=<virtual disk name>

To view the list of configured virtual disks, use the command **SHOW VIRTUAL_DISK *** (Figure 51).

Figure 51. Show Virtual Disk Information Screen

```
RAID[0]$ show virtual_disk *

                                                                    |      Home       |Background|
Idx|Name           | State |Pool|Raid|Cap GB|Settings| Jobs  |Current|Preferred|   Job    |
-------------------------------------------------------------------------------------------------
  0 vd-0_0          READY     0  5     3632   W   I          0(L) 0   0(L) 0   INACTIVE
  1 vd-1_1          READY     1  5     3632   W   I          0(L) 0   1(R) 0   INACTIVE
  2 vd-2_2          READY     2  5     3632   W   I          0(L) 0   0(L) 0   INACTIVE
  3 vd-3_3        NOT RDY     3  5     5480   W   I   I      0(L) 0   1(R) 0   0x28000003

Total Virtual Disks: 4
```

To display the detailed information of a virtual disk, use the command:
SHOW VIRTUAL_DISK=<id> ALL_ATTRIBUTES (Figure 52).

Figure 52. Show Virtual_Disk All_Attributes Example Screen

```
**************************
*     Virtual Disk(s)     *
**************************
  Index:            0
  OID:              0x8c7b0000
  Name:             vd-0_0
  Pool Index:       0
  Pool OID:         0x1c790000
  Capacity:         16 GB
  Offset:           0x0
  State:            READY
  Raid Level:       RAID6
  IO ROUTING:       TRUE
  WBC:              TRUE
  MIRRORED:         TRUE
  Initializing:     FALSE
  Paused:           FALSE
  AutoWriteLock:    FALSE
  Critical:         FALSE
  Forced Write-thru:FALSE
  Current Home:     0x38000000  0x00000000 (REMOTE)
  Future Home:      0xffffffff  0x00000000
  Preferred Home:   0x38000000  0x00000000 (REMOTE)
  Job OID:          INACTIVE
  UUID:             60001ff08025a000000000008c7b0000
```

2.7.11 Create and Assign Spare Pools

The SFA10000 supports the concept of spare pool. A spare pool contains physical disks that can be used as spare disks in one or more storage pools.

In the event of disk failure, a disk that is assigned to a spare pool is automatically swapped into the storage pool to replace the failed disk. A rebuild then occurs immediately, reducing the possibility of data loss due to additional disk failures.

Spare pools can be shared by storage pools or dedicated to a particular storage pool in which case the spare disks will only be used by the designated storage pool. After a spare pool is created, you will need to select disk(s) from the Unassigned Pool and assign it to the spare pool.

Follow these steps to create a spare pool:

1. At the CLUI prompt, enter the command: **CREATE SPARE_POOL** (Figure 53)

   A message is displayed to indicate whether the new spare pool creation was successful.

   Figure 53. Create Spare Pool Example Screen

   ```
   RAID[0]$ create spare_pool
   SPARE POOL 6 OID=0x19bc0006 create STATUS='Success' (0x0)
   ```

2. Assign physical disk(s) to the spare pool. Enter the command:
   ASSIGN PHYSICAL_DISK=<disk-id> TO_POOL=<pool-id>
   where <disk-id> is the index name of the disk to be added and <pool-id> is the index name of the spare pool.

   A message is displayed to indicate whether the disk has been added to the spare pool (Figure 54).

   Figure 54. Assign Physical Disk to Spare Pool Example Screen

   ```
   RAID[0]$ assign physical_disk=648 to_pool=6
   PHYSICAL_DISK 648 OID=0x22b10288 assigned to POOL 6 OID=0x19bc0006STATUS='Success' (0x0)
   ```

   To view the list of available disks, enter the command:
   SHOW UNASSIGNED_POOL PHYSICAL_DISK

3. Use the **CREATE SPARE_POOL** and **ASSIGN PHYSICAL_DISK** commands to create more spare pools as needed.

   If you need to delete a disk from the spare pool, enter the command:
   ASSIGN PHYSICAL_DISK=<disk-id> TO_POOL=0x1800ffff

   To delete a spare pool, enter the command:
   DELETE SPARE_POOL=<pool-id>

4. You may assign a name to a spare pool using the command:
   SET SPARE=<spare pool-id> NAME=<name>

5. To view the list of configured spare pool(s), enter the command:
   SHOW SPARE_POOL * (Figure 55).

Figure 55. Show Spare Pool Example Screen (1)

```
RAID$ show spare_pool *
                              |Disk|Total |Total|Storage|
Idx|Name             |Blocks|T/O|Cap GB| PDs |Pool Idx|
-------------------------------------------------------
  6 spare_pool-1       512   10    120    1
Total Spare Pools: 1
```

To display the detailed information of a spare pool (Figure 56), use the command:
SHOW SPARE_POOL=<id> ALL_ATTRIBUTES

To display the list of disks in a spare pool (Figure 56), use the command:
SHOW SPARE_POOL=<id> PHYSICAL_DISKS

Figure 56. Show Spare Pool Example Screen (2)

```
RAID$ show spare_pool=6 all_attributes
  Index:               6
  OID:                 0x19bc0006
  Type:                GLOBAL SPARE
  Name:                spare_pool-1
  Block Size:          512
  DiskTimeout(FRT):    10 minutes
  Total Capacity:      120 GB
  UUID:                0x0000000000000000
  Total Phy Disks      1
  Storage Pool List:
Total Spare Pools: 1

RAID$ show spare_pool=6 physical_disks
                                                    |Health|
Encl|Slot|Vendor|Product ID| ::::::::::|Pool|State | Idx|State|WWN|
-------------------------------------------------------------------
   2   53 HITACHI : : : : : : : : : :    6 SPARE   648 NORM  :::
```

6. Assign spare pools to storage pools. Enter the command:
SET POOL=<pool-id> SPARE_POOL=<spare-pool-id>
where <pool-id> and <spare-pool-id> are the index names of the storage pool and spare pool respectively. Repeat this step for each storage pool (Figure 57).

Figure 57. Set Storage Pool to Spare Pool Example Screen

```
RAID[0]$ set pool 0 spare_pool 6
POOL 0 OID=0x19b40000 attributes set with STATUS='Success' (0x0)
```

NOTE : Each storage pool should have a spare pool assigned to it. If a storage pool's attributes show "Global Spare Pool: UNASSIGNED", no spare pool will serve this storage pool.

Use the **SHOW POOL** command to view the new storage pool information (Figure 58).

Figure 58. Show Pool Example Screen

```
RAID$ show pool=1 all
   Index:               0
   OID:                 0x180d0000
   Type:                STORAGE
   Name:                pool-0
    :
    :
   Total Capacity:      600 GB
   UUID:                60000000000000000000180d0000
   Global Spare Pool:   6
   DiskTimeout(FRT):    10 minutes
    :
    :
```

2.7.12 Present Virtual Disk to External Host (SFA10000 Only)

NOTE : For SFA10000E, this feature is not applicable.

Virtual disks are only presented to the hosts that have been given authorized access. A Presentation on a SFA10000 has the following components:

- Discovered Initiators
  A Discovered Initiator is defined as a Fibre Channel or InfiniBand initiator that has logged into the SFA10000. This is normally a single port on a Fibre Channel or InfiniBand HBA in an external computer.

- Host
  A Host is an abstract client of a virtual disk that has all the access to all data within the virtual disk. This most commonly represents an external computer with one or more Fibre Channel or InfiniBand initiators. Host attributes are OS (operating system) of which you can select GENERIC, WINDOWS, LINUX or MAC_OS. The OS attribute dictates the way the controller responds to certain Fibre Channel commands where appropriate. For example, Mac OS X expects a different response from an inquiry of LUN0 (the controller LUN) than a Windows host.

- Channel
  A Channel is one of the Fibre Channel or InfiniBand ports on the controller. Channel attributes are MODE of which you can select either MAC_OS or STANDARD. When MAC_OS is selected, the port's node name will be set differently in order to be visible to a Macintosh system.

- Stack
  A Stack is the input side of the subsystem. In the case of the SFA10000, the stack is defined as the Fibre Channel interface. In future versions of the product, different interface stacks may exist.

- Presentation
  A Presentation is the relationship between a Host and a virtual disk. A Presentation implies that the related Host has some sort of access to the virtual disk. Attributes of a Presentation are:
  - PORT: from which the host will see the virtual disk
  - READ_ONLY: controls read only access
  - PRESENT_HOME_ONLY: presents the specified virtual disk from its designated home controller only
  - LUN: user-specified LUN number that the virtual disk will show to the host.

Follow these steps to set up a presentation:

1. To display the currently available initiators, enter the command (Figure 59):
APPLICATION SHOW DISCOVERED_INITIATOR

Figure 59. Display Available Initiators

```
RAID$ app show discovered_initiator

                |         Initiator Identifier          | Controller 0  | Controller 1  |
Index |Type|   ID   |    node           |      port         | RP 0  | RP 1  | RP 0  | RP 1  |
-------------------------------------------------------------------------------------------
00001   FC |0x0000e8 0x2001001b32ae096c 0x2101001b32ae096c          |       |     3 |
00002   FC |0x0000e8 0x2001001b32aeb580 0x2101001b32aeb580          |       |       |0
00003   FC |0x0000e8 0x2001001b328e0280 0x2001001b328e0280          |       |    2  |
00004   FC |0x0000e8 0x2001001b32ae176c 0x2001001b32ae176c          |       |       |  1
00005   FC |0x0000e8 0x2001001b328eb580 0x2001001b328eb580          |0      |       |
00006   FC |0x0000e8 0x2001001b32aeb280 0x2001001b32aeb280       2  |       |       |
00007   FC |0x0000e8 0x2001001b328e176c 0x2001001b328e176c          |  1    |       |
00008   FC |0x0000e8 0x2001001b328e096c 0x2101001b328e096c        3|       |       |

Total FC Initiators: 8
```

2. Create a host. Enter the command (Figure 60):
APPLICATION CREATE HOST NAME=<host name>
OSTYPE=[LINUX|WINDOWS|MAC_OS|DEFAULT|GENERIC]
where <host name> is an assigned host name to help make mapping simpler for the user; <os type> is the mode which can be set to characteristics specific to an Operating System, especially for Mac OSX.

   Use the **APPLICATION CREATE HOST** command to create more hosts as needed.

   If you need to delete a host, use the **APPLICATION DELETE HOST=<host-id>** command.

   To display the current settings, enter the command: **APP SHOW HOST**

Figure 60. Create Host Examples

```
RAID$ app create host name=server1-port1 ostype=linux
HOST 0 OID=0x18d0000 creation STATUS='Success' (0x0)

RAID$ app create host name=server1-port2 ostype=linux
HOST 1 OID=0x18e0001 creation STATUS='Success' (0x0)

RAID$ app show host
                          |Stack| Host Mode Attributes  |
Index|Host Name            |Index|OS Type|Characteristics|
-------------------------------------------------------------
00000 server1-port1         00000 LINUX   0x0000000000000001
00001 server1-port2         00001 LINUX   0x0000000000000001

Total Hosts: 2
```

3. Map a host to a discovered initiator. Enter the command (Figure 61):
APPLICATION IMPORT DISCOVERED_INITIATOR=<initiator-id>
HOST=<host-id>
where <initiator-id> is the index name of the discovered initiator; <host-id> is the index name of the host.

Figure 61. Map a Host to a Discovered Initiator

```
RAID$ app import discovered_initiator=6 host=0

INITIATOR 0 OID=0x280f0000 imported from DISCOVERED_INITIATOR 6 oid=0X30000006
STATUS='Success' (0x0)
```

Use the **APPLICATION IMPORT** command to map the other hosts to the initiators as needed.

To display the current mappings, enter the command (Figure 62):
APPLICATION SHOW INITIATOR

Figure 62. Show Current Imported Initiators

```
RAID$ app show initiator

           |Host |        Initiator Identifier          |
Index |Type|Index|      node         |      port        |
-------------------------------------------------------------
00000   FC  00000 0x2001001b32aeb280 0x2001001b32aeb280

Total FC Initiators: 1
```

4. Present the virtual disks to the hosts. Enter the command (Figure 63):
APPLICATION CREATE PRESENTATION HOST=<host-id> VIRTUAL_DISK=<vd-id> LUN=<LUN-id>
where <host-id> is the index name of the host; <vd-id> is the virtual disk to be presented; <LUN-id> is the LUN that the specified virtual disk will be presented as or use the default value if it is not specified.

Figure 63. Create a Presentation

```
RAID$ app create presentation host=0 vd=0 lun=0

PRESENTATION 0 OID=0x0x20110000 creation STATUS='Success' (0x0)
```

To simply present the virtual disk to all host ports for both controllers (promiscuous mode), use this command:
APPLICATION CREATE PRESENTATION VIRTUAL_DISK=<vd-id> HOST=ALL

Use the **APPLICATION CREATE PRESENTATION** command to configure other presentations as needed.

To display the current settings, enter the command (Figure 64):
APPLICATION SHOW PRESENTATION *

Figure 64. Show Current Presentations

```
RAID$ app show presentation *
                                                  |      Channel Mask         |
                                                  | Controller 0|Controller 1 |
Pres.|Host           Host | VD  |     |Home|Read|  RP 0| RP 1 |  RP 0| RP 1 |
Index|Name           Index|Index| LUN |Only|Only| 0123 | 0123 | 0123 | 0123 |
-------------------------------------------------------------------------------
00000 server1-port1   00000 00000  000  OFF  R/W    f      f      f      f

Total Presentations: 1
```

2.7.12.1 Special Considerations for MAC OS®

Apple's Mac OS X Server does not adhere to the FC specification with regards to node naming. Apple expects the node name to vary from port to port while the FC specification calls for it to remain constant. A mode setting is provided to force the node name to be unique on a given client channel (Fibre Channel port). This mode can be set for one or more client channels. While the feature is aimed at Mac SANs, it can also be used for connections to heterogeneous client SANs.

When creating the host, specify the **OSTYPE** to be MAC_OS:
APPLICATION CREATE HOST OSTYPE=MAC_OS

Refer to Section 3.6.4 "Additional Configuration Considerations for Mac® Hosts" for more information.

2.7.12.2 Selective Presentation

By default, a LUN is accessible via all initiators/ports. You may mask a presentation and choose the specific port on which the initiator may have access to the LUN.

You may mask a presentation when you create it:
APPLICATION SET PRESENTATION HOST=<host-id> VIRTUAL_DISK=<vd-id> LUN=<LUN-id> ENABLE=<mask option>

By selecting the **ENABLE** option, you are choosing the specific port on which the initiator may have access:

Controller 0 RP0: **ENABLE=0**
Controller 0 RP1: **ENABLE=1**
Controller 1 RP0: **ENABLE=2**
Controller 1 RP1: **ENABLE=3**

To mask a LUN to all ports, use **ENABLE=ALL**.

To change a mask, you must first reset it with **ENABLE=NONE**.

To mask a LUN to multiple ports, you must set one port at a time. For example, to mask the same LUN 4 to Controller 0, RP0 and Controller 1, RP1:

APP SET PRESENTATION=4 ENABLE=0
APP SET PRESENTATION=4 ENABLE=3

2.7.13 Storage Pool Initialization

When a storage pool is created, initialization begins automatically and will continue in the background until it is completed. While its virtual disk(s) can be accessed immediately, the performance of the virtual disks will be degraded while it continues to initialize.

To monitor the progress of a job, use the **SHOW JOB * ALL_ATTRIBUTES** command. It will display the type of job and the percentage of completion (Figure 65).

Figure 65. Show Job Example Screen

```
RAID[0]$ show job * all_attributes
  OID:                  0x2e040001 (Index:1)
  Target:               POOL:0x1aa00001 (Index:1)
  Sub-Target:           NA
  Type:                 INITIALIZE
  State:                RUNNING
  Completion Status:UNKNOWN
  Priority:             50
  Fraction Complete:71%
  Time:                 NA
```

You may also check if a virtual disk is ready for access using the **SHOW VIRTUAL_DISK <id> ALL_ATTRIBUTES** command (Figure 66).

Figure 66. Show Virtual Disk Example Screen

```
RAID[0]$ show virtual_disk 0 all_attributes
  OID:                  0x89ba0000
  Name:                 89ba0000
  POOL OID:             0x19b40003
  Capacity:             16384 MBs
  Offset:               0x0
  State:                READY
  Raidlevel:            RAID5
  IO ROUTING:           TRUE
  WBC:                  TRUE
  MWBC:                 FALSE
  Initializing:         FALSE
  Paused:               FALSE
  AutoWriteLock:        FALSE
  Data Lost:            FALSE
  Present Home Only:FALSE
   .
   .
```

Chapter 3

Administration

3.1 Managing the SFA10000

The SFA OS provides a set of tools that enable administrators to centrally manage the network storage and resources that handle business-critical data. These include Configuration Management, Performance Management, and Firmware Update Management.

NOTE : The command descriptions given in this User Guide are specific to Firmware Version 1.4.0. To access the most up-to-date commands, use the CLUI's Online Help feature.

3.1.1 Management Interface

SAN management information for the SFA10000 can be accessed locally through a serial interface, or remotely through SSH.

Locally via Serial Interface

Any RS-232 terminal or terminal emulator (such as Hyperterminal) can be used to configure and monitor the SFA10000.

1. Connect the supplied serial cable to the RS-232 port on the controller (Figure 67). Connect the other end of the cable to a serial port on a standard PC.

Figure 67. RS-232 Port on Controller

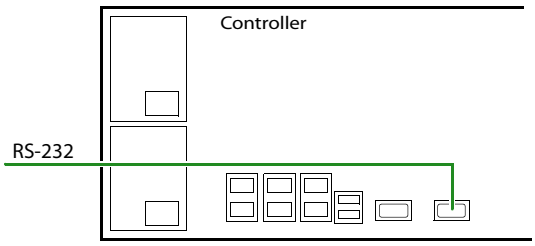

2. Load a serial console program (such as HyperTerminal, minicom, and screen) and use the following settings for the serial connection:
   - Bits per second: 115,200
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: None

3. Once connected, press the <Enter> key to bring up the login prompt. Enter the user name **`user`** and password **`user`**.

Remotely via SSH

To configure and monitor the SFA10000 remotely, connect the system to your Ethernet network. Please refer to Section 3.13, "Remote Management of SFA10000" for further information on how to set up the controller's network interface.

3.1.2 User Logins

RS-232 Login

For a terminal session, press the <Enter> key to bring up the login prompt. Enter the user name **user** and password **user**.

Login via SSH

For a SSH session (Figure 68), enter the user name **user** and password **user**. Both are case sensitive.

Figure 68. SSH Login Screen

```
login as: user
user@10.23.23.16's password:

Linux (none) 2.6.25-ddn-016620-3 #2 SMP Mon Aug 16 10:38:29 MST 2010 x
DataDirect Networks SFA
```

Logout

To logout, enter the command **QUIT**. For SSH connection, the current session will be disconnected.

3.1.3 Available Commands

The SFA OS CLUI commands are independent of case. Most of the keywords can be abbreviated and most of the punctuations are optional.

Entering the "**?**" character at the CLUI prompt will display the available commands within the CLUI (Figure 69). To get help information on a command, type the command followed by a question mark.

For example, **ASSIGN ?**<Enter> will display help on "assign" options on the system.

Figure 69. Help Screen

```
RAID[0]$ ?
Select one of the following subjects for additional commands:
   APPLICATION      Selects the application processor as the subject
   RAID             Selects the raid processor as the subject
   UI               Selects the user interface as the subject

Additional commands:
   HELP             Provides information on how to use the user interface help
   QUIT             Exits the Command Line User Interface

Object class options available for default subject RAID:
Available keywords:
   ASSIGN           Assign raid object
   CLEAR            Clear raid object states
   CREATE           Create raid objects
   DELETE           Delete raid objects
   ERASE            Erase raid objects
   LOCATE           Locate raid objects
   Mark             Mark the event log
   MOVE_HOME        Locate raid objects
   PAUSE            Pause raid objects
   REPLACE          Replace raid objects
   RESUME           Resume raid objects
   SET              Set raid objects
   SHOW             Display raid objects
   SHUTDOWN         Shutdown raid objects
   SYNCHRONIZE      Synchronize raid objects
   UPDATE_FIRMWARE  Update firmware on raid objects
   VERIFY           Start a Verify background job(s)
```

3.1.3.1 Basic Key Operations

The command line editing and history features support **ANSI** and **VT-100** terminal modes. The command history buffer can hold up to 64 commands. The full command line editing and history only work on main CLI and SSH sessions when entering new commands. Basic Key Assignments are listed in Table 2.

Simple, not full command, line editing only is supported when the:

- CLUI prompts the user for more information.
- alternate CLUI prompt is active. (The alternate CLUI is used on the RS-232 connection during an active SSH session.)

NOTE : Not all SSH programs support all the keys listed in Table 2.

Table 2. Basic Key Assignments

Key	ANSI CTRL or Escape Sequence	Description
Backspace	Ctrl-H	deletes preceding character
Delete	Del, or Esc [3~	deletes current character
Up Arrow	Esc [A	retrieves previous command in the history buffer
Down Arrow	Esc [B	retrieves next command in the history buffer
Right Arrow	Esc [C	moves cursor to the right by one character
Left Arrow	Esc [D	moves cursor to the left by one character
Home	Esc [H or Esc [1~	get the oldest command in the history buffer
End	Esc [K or Esc [4~	get the latest command in the history buffer
Insert	Esc [2~	toggles between insert mode and overtype mode
PgUp	Esc [5~	retrieves oldest command in the history buffer
PgDn	Esc [6~	retrieves latest command in the history buffer
Ctrl-U	Ctrl-U	delete to beginning of line
Ctrl-K	Ctrl-K	delete to end of line
Ctrl-A	Ctrl-A	move cursor to beginning of line
Ctrl-E	Ctrl-E	move cursor to end of line

3.2 Physical Disk Information

The disks in the system are categorized into one of the following pools:

- Unassigned Pool - By default, all the disks or replacement disks are initially allocated to this pool. When a disk has failed or been released (from a spare pool, for example), it is also moved to this pool.
- Storage Pool (sometimes referred to as a RAID group) - A collection of 2 to 10 disks, ideally of the same capacity and type.
- Spare Pool - This type of pool contains Physical Disks that can be used as spare disks in one or more storage pools.

3.2.1 Disk Information

To display the list of all the disks installed in the system, enter the command (Figure 70):

SHOW PHYSICAL_DISK *

Figure 70. Show Physical Disk List Example Screen

```
**********************
*   Physical Disks   *
**********************
Encl|Slot| Vendor |     Product ID     |Type|Cap GB | RPM|Revision|    Serial Number   |Pool|State | Idx|State |       WWN       |
-----------------------------------------------------------------------------------------------------------------------------------
   1    1 SEAGATE  ST32000444SS          SAS    1863 7.2K 0006     9WM0DRN00000C0331WFJ    0  GOOD   139 NORM    5000c50010444e40
   1    2 SEAGATE  ST32000444SS          SAS    1863 7.2K 0006     9WM0HBM10000C034D008    0  GOOD   138 NORM    5000c500104a1370
   1    3 SEAGATE  ST32000444SS          SAS    1863 7.2K 0006     9WM0FY4G0000C0372EJE    0  GOOD   128 NORM    5000c50020c7c8f0
   1    4 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK043HW0000C0283GKU UNAS  GOOD   190 READY   5000c500103939fc
   1    5 SEAGATE  ST32000444SS          SAS    1863 7.2K 0006     9WM0JVDG0000C032G0Q1    0  GOOD   180 NORM    5000c50020c43740
   1    6 SEAGATE  ST32000444SS          SAS    1863 7.2K 0006     9WM0JZ510000C034CXFP    0  GOOD   177 NORM    5000c50020c40f5c
   1    7 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK08SGR0000C035AQK0 UNAS  GOOD   199 READY   5000c50020cb02a4
   1    8 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06ACN0000C034A5TG UNAS  GOOD   174 READY   5000c500104ec228
   1    9 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06EMR0000C032G0VJ UNAS  GOOD   192 READY   5000c500104f9d84
   1   10 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK07KJ20000C032B05W    2  GOOD   130 NORM    5000c50020caf5e0
   1   11 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06BXY0000S005SDYV UNAS  GOOD   170 READY   5000c500104f7e5c
   1   12 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06SNV0000C0268B7K UNAS  GOOD   193 READY   5000c500104ded34
   1   13 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK08HZC0000C035HZ5Z UNAS  GOOD   166 READY   5000c50020cb1e20
   1   14 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06C4N0000C0325YY2 UNAS  GOOD   169 READY   5000c500104f8550
   1   15 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK08JEB0000S018YUL0 UNAS  GOOD   179 READY   5000c50020cb1158
   1   16 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06C3F0000C0326TSF UNAS  GOOD   178 READY   5000c500104f7b78
   1   17 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK068QL0000C035J0U4 UNAS  GOOD   153 READY   5000c500104f1c94
   1   18 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK089DZ0000C029F1JW UNAS  GOOD   186 READY   5000c50020cb057c
   1   19 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06XHV0000C0212ZAD UNAS  GOOD   230 READY   5000c50020c66c9c
   1   20 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK08EB40000C0280TT5    2  GOOD   124 NORM    5000c50020caf780
   1   21 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK071T70000C035HZX5 UNAS  GOOD   194 READY   5000c500104f73ac
   1   22 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK07RN20000C0372K9N UNAS  GOOD   146 READY   5000c50020cb0414
   1   23 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK081ZD0000C0372KB9 UNAS  GOOD   200 READY   5000c50020cb2388
   1   24 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK05P0G0000C0317LG8    2  GOOD   137 NORM    5000c500104c7600
   1   25 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06WNG0000C0365LAD UNAS  GOOD   145 READY   5000c50020c61504
   1   26 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK069SL0000C031ANA4 UNAS  GOOD   210 READY   5000c500104ebf98
   1   27 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK08HTL0000C032AYSB UNAS  GOOD   208 READY   5000c50020cb16ac
   1   28 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK08EB00000C032H8XE UNAS  GOOD   224 READY   5000c50020cb0bec
   1   29 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06QY40000C0368V2V UNAS  GOOD   235 READY   5000c50020c5fd24
   1   30 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK0613T0000C036BKJH UNAS  GOOD   144 READY   5000c50020c60d0c
   1   31 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06BGZ0000C036420K UNAS  GOOD   154 READY   5000c500104f7c34
   1   32 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK0621D0000C036FLVE UNAS  GOOD   183 READY   5000c50020c60974
   1   33 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK08APH0000C035HZNU UNAS  GOOD   196 READY   5000c50020cb2a70
   1   34 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK08SFH0000C036FMKT UNAS  GOOD   150 READY   5000c50020cb031c
   1   35 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK07QJX0000C0368UBU UNAS  GOOD   231 READY   5000c50020c69d6c
   1   36 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK03B3A0000S018WEQS UNAS  GOOD   191 READY   5000c50010334cf0
   1   37 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK08RPJ0000C035D4F7    2  GOOD   131 NORM    5000c50020cafaf4
   1   38 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK03Q4Z0000C03008S1 UNAS  GOOD   160 READY   5000c50010386b68
   1   39 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK048GN0000C0358PJD UNAS  GOOD   195 READY   5000c50020cb0990
   1   40 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK04XPX0000C030JNQV UNAS  GOOD   167 READY   5000c50010468258
   1   41 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06X410000C031DRBQ UNAS  GOOD   226 READY   5000c500104fdfa4
   1   42 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK07XAT0000C0214TUW UNAS  GOOD   155 READY   5000c50020cb0e18
   1   43 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK08RZK0000C030EKKL    2  GOOD   129 NORM    5000c50020cafdcc
   1   44 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06PDX0000C032H8PR UNAS  GOOD   149 READY   5000c500104c07a4
   1   45 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06AWC0000C029E82U UNAS  GOOD   216 READY   5000c500104ee58c
   1   46 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK0763Z0000C03611SH UNAS  GOOD   185 READY   5000c500104f6ca0
   1   47 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK075RV0000S006TK08 UNAS  GOOD   182 READY   5000c500104f62a0
   1   48 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06LRG0000C0294GYC UNAS  GOOD   168 READY   5000c50020c60248
   1   49 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK04SMH0000C0214THP UNAS  GOOD   171 READY   5000c50010486f74
   1   50 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06NMJ0000C0358Q71 UNAS  GOOD   211 READY   5000c500104c9ac0
   1   51 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK06WHH0000C03190PT UNAS  GOOD   215 READY   5000c50020c62dbc
   1   52 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK05Y7P0000C0366HN6 UNAS  GOOD   225 READY   5000c50020c62dd0
   1   53 SEAGATE  ST31000424SS          SAS     931 7.2K 0006     9WK0772L0000C0368V85 UNAS  GOOD   162 READY   5000c50020c62518
    :
    :

Total Physical Disks:       300
Total Assigned Disks:       22
Total Unassigned Disks:     98
  Total SAS Disks:          300
  Total Member State Normal:   22
```

To display the list of unassigned disks, enter the command:
SHOW UNASSIGNED_POOL PHYSICAL_DISK

To retrieve information about a specific disk, enter the command:
SHOW PHYSICAL_DISK=<disk-id> ALL_ATTRIBUTES (Figure 71).

Figure 71. Physical Disk Information Example Screen

```
**************************
*     Physical Disk(s)     *
**************************
  Index:                852
  OID:                  0x20770354
  Pool Index:           UNASSIGNED
  Pool OID:             UNASSIGNED
  Capacity:             536 GB
  Raw Capacity:         558 GB(Base 2)/600 GB(Base 10)
  Block Size:           512
  Available Channels:   02  NA  NA  NA  NA  NA  NA  NA  14  NA  NA  NA  NA  NA  NA  NA
                        22  NA  NA  NA  NA  NA  NA  NA  36  NA  NA  NA  NA  NA  NA  NA
  Enabled Path IDs:     0x174  0x20d0  0x218  0x2218
  Enclosure Index:      1
  Enclosure OID:        0x50000001
  Disk Slot:            1  (1:1)
  Vendor ID:            HITACHI
  Product ID:           HUS156060VLS600
  Product Revision:     A5D0
  Serial Number:        JWV3LUPE
  Health State:         GOOD
  Rotation Speed:       15000 RPM
  Device Type:          SAS
  Member State:         UNASSIGNED
  State:                READY
  Spare:                FALSE
  Failed:               FALSE
  UUID:                 0x5000cca00f068f2c
  Interposer Type:      NONE
  Update in Progress:   FALSE
```

If there is a failed disk, use the **SHOW UNASSIGNED_POOL FAILED ALL** command to display the failed disk's information (Figure 72).

Figure 72. Show Failed Disk Example Screen

```
**************************
*     Physical Disk(s)     *
**************************
  Index:              139
  OID:                0x2135008b
  Pool Index:         UNASSIGNED
  Pool OID:           UNASSIGNED
  Capacity:           1824 GB
  Raw Capacity:       1863 GB(Base 2)/2000 GB(Base 10)
  Block Size:         512
  Enabled Disk Ch:    0x2b4  0xd1
  Enclosure Index:    1
  Enclosure OID:      0x50000001
  Disk Slot:          1  (1:1)
  Vendor ID:          SEAGATE
  Product ID:         ST32000444SS
  Product Revision:   0006
  Serial Number:      9WM0DRN00000C0331WFJ
  Health State:       FAILED
  Rotation Speed:     7200 RPM
  Device Type:        SAS
  Member State:       UNASSIGNED
  State:              READY
  Spare:              FALSE
  Failed:             TRUE
  UUID:               0x5000c50010444e40
  Update in Progress:   FALSE
```

3.2.2 Disk States

Listed below are the possible disk states:

GOOD: Disk is in a normal or functional condition.

AMIS: Disk is Already MISsing.
The disk in this state must return prior to the other disks. If the specific disk that is in this state never returns to the pool, the pool will remain inoperative. Replacing this disk to do a rebuild will do nothing.

WTRB: Disk is Waiting To ReBuild.
A disk went missing then came back and is rebuilding (either partial or full). Then other disks in the storage pool went missing causing it to go inoperative, which halts the rebuild. Now waiting for other disks in the system to become active again.

MNRB: Disk is Missing with No ReBuild.
This is the same as missing, but when the disk comes back, there is nothing to rebuild. This could happen when multiple disks fail at the same time making the storage pool inoperative. Reseating is an option here to see if the disk would come back. This is similar to the AMIS condition, but in this case it does not matter what order the MNRB disks become active.

MISS: Disk is MISSing.
This is either not electrically active or seen, or only one controller sees it in a couplet system. Reseating would be a first recommendation as a partial rebuild would happen if it then appears to both controllers. Replacing would cause a full rebuild.

FAIL: Disk is FAILed and is unassigned.
This means that the disk timeout has either expired or the disk had errors and had been failed. "FAIL" in the storage pool command means that it has been failed out of the storage pool, not necessarily the state of the disk. This is being renamed to EMPTY.

RBLD: Disk is ReBuiLDing.

PRTRDY: PaRTial ReaDY.
This state is a condition where only one of the two controllers is able to communicate with a drive.

3.2.3 Visual Indication

`LOCATE PHYSICAL_DISK=<id>` provides a visual indication of the specified disk. The status LED of the disk module will blink.

`LOCATE UNASSIGNED_POOL` provides a visual indication of the disk modules that are unassigned. The status LED of the disk modules will blink.

`LOCATE UNASSIGNED_POOL FAILED` provides a visual indication of the disk modules that have failed. The status LED of the disk modules will blink.

3.3 Storage Pool Management

The SFA10000 creates centrally-managed and vendor-independent storage pooling. It enables different types of storage to be aggregated into a single logical storage resource from which virtual volumes (virtual disks) can be served up to multi-vendor host computers. The networked storage pools will provide the framework to manage the growth in storage demand from web-based applications, database growth, network data-intensive applications, and disaster tolerance capabilities.

3.3.1 Display Storage Pool Information

You can add and remove storage pools without affecting system operations. Use the **SHOW POOL *** command to display the current list of storage pools (Figure 73).

Figure 73. Show Storage Pool List Example Screen

```
*****************
*    Pool(s)    *
*****************                                        |Total |Free  |Max   |                  |Disk|  Global  |Spare |
Idx|Name           |State      |Chunk|Raid| Faults |cap GB|cap GB|VD GB | Settings  | Jobs  |T/O |spare pool|Policy|
-------------------------------------------------------------------------------------------------------------------------
  0 pool-0         NORMAL       128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO
  1 pool-1         NORMAL       128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO
  2 pool-2         NORMAL       128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO
  3 pool-3         NORMAL       128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO
  4 pool-4         NORMAL       128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO
  5 pool-5         NORMAL       128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO

Total Storage Pools: 5
```

NOTE : If you issue the command, **SHOW POOLS ***, and the pools indicate there is a fault, check the details of the pool. The pool will go into WR thru state on initial startup to allow the batteries the opportunity to charge. The WR Thru mode is necessary to allow for data integrity until battery redundancy on each controller is completed. If there is a battery issue, the system will indicate replacement is required and all VDs will remain in WR Thru mode until the issue is resolved.

The **SHOW POOL=<id> ALL_ATTRIBUTES** command displays the detailed information of the specified storage pool (Figure 74).

Figure 74. Show Storage Pool Attributes Example Screen

```
******************
*     Pool(s)     *
******************
  Index:                0
  OID:                  0x1c790000
  Type:                 STORAGE
  Name:                 pool-0
  Chunk Size:           128KB (0x100 blocks)
  Block Size:           512
  RAID Type:            RAID6
  Free Raid6 Capacity:  21784 GB
  Max VD Size:          21784 GB
  Total Capacity:       27360 GB
  UUID:                 60001ff08025a000000000001c790000
  Global Spare Pool:    0x18d60013  (Index 19)
  DiskTimeout(FRT):     10 minutes
  Init Policy:          ALLOW_IO
  Init Priority:        50%
  Verify Priority:      70%
  Full Rebuild Priority:        80%
  Fractional Rebuild Priority:  90%
  Sparing Policy:       AUTOMATIC
  Verify Policy:        ENABLED
  Assign Policy:
    Device Type:        NA
    Rotation Speed:     NA
    Raw Capacity:       2794 GB
  SATAssure:            None
  Cache Settings:
    ReACT:              FALSE
    IO Routing:         TRUE
    Mirroring:          TRUE
    Read Ahead:         TRUE
    Write Back:         TRUE
  Initializing:         FALSE
  Rebuilding:           FALSE
  Paused:               FALSE
  AutoWriteLock:        FALSE
  Critical:             FALSE
  Forced Write-Thru     FALSE
  Current Home:         0x38000000  0x00000000 (REMOTE)
  Future Home:          0xffffffff  0x00000000
  Preferred Home:       0x38000000  0x00000000 (REMOTE)
  Job OID:              INACTIVE
  Total Phy Disks       10
  State:                NORMAL
  Member Size:          2736 GB
    pID     State     UUID
    0x0459  NORM      0x5000cca225c2b913
    0x0329  NORM      0x5000cca225c2c89c
    0x0462  NORM      0x5000cca225c2a53f
    0x0466  NORM      0x5000cca225c2b738
    0x046c  NORM      0x5000cca225c2a653
    0x0470  NORM      0x5000cca225c2b179
    0x030c  NORM      0x5000cca225c2d1ba
    0x046f  NORM      0x5000cca225c28b8d
    0x0479  NORM      0x5000cca225c2dadf
    0x0473  NORM      0x5000cca225c28bce
Total Storage Pools: 1
```

The **SHOW POOL=<id> PHYSICAL_DISK** command displays the list of disks associated with the specified storage pool (Figure 75).

Figure 75. Show Storage Pool Physical Disks Example Screen

```
**************************
*     Physical Disk(s)    *
**************************
                                                                                                                |Health|
Encl|Slot| Vendor |     Product ID     |Type|Cap GB | RPM|Revision|    Serial Number   |Pool|State | Idx|State |       WWN       |
-------------------------------------------------------------------------------------------------------------------------------------
   1   37 HITACHI  HUS153014VLS300       SAS     136  15K A410     JFVWMZNC5301           0  GOOD   611 NORM    5000cca00d324438
   1   39 HITACHI  HUS153014VLS300       SAS     136  15K A140     J4V1YYAA5301           0  GOOD   481 NORM    5000cca005039324
   1   51 HITACHI  HUS153014VLS300       SAS     136  15K A410     JFVWL5YC5301           0  GOOD   479 NORM    5000cca00d32293c
   2    1 HITACHI  HUS153014VLS300       SAS     136  15K A410     JFVSU5PC5301           0  GOOD   483 NORM    5000cca00d2b4c48
   2    3 HITACHI  HUS153014VLS300       SAS     136  15K A410     JFV8XYGC5301           0  GOOD   487 NORM    5000cca00d103e94
   2   42 HITACHI  HUS153014VLS300       SAS     136  15K A410     JFVPKHUC5301           0  GOOD   477 NORM    5000cca00d273618

Total Physical Disks:      6
Total Assigned Disks:      6
Total Unassigned Disks:    0
  Total SAS Disks:         6
  Total Member State Normal:   6
```

The **SHOW POOL=<id> VIRTUAL_DISK** command displays the list of virtual disks associated with the specified storage pool (Figure 76).

Figure 76. Show Storage Pool Virtual Disks Example Screen

```
RAID[0]$ show pool 0 vd

                                                                    |          Home    |Background|
Idx|Name            | State  |Pool|Raid|Cap GB|Settings| Jobs   |Current|Preferred|    Job     |
-------------------------------------------------------------------------------------------------
  0 vd-0_0           READY     0  5     3632    W    I            0(L) 0    0(L) 0    INACTIVE

Total Virtual Disks: 1
```

3.3.2 Creating a Storage Pool

A storage pool on the SFA10000 has the following attributes:

- RAID Level
  Storage pools can be configured to use either a RAID 1, RAID 5, or RAID 6 parity scheme. In RAID 1, the capacity of one disk is used for data duplication. In RAID 5, the capacity of one disk is reserved for parity, allowing data recovery in the event of a single disk loss in the storage pool. In RAID 6, the capacity of two disks is reserved for parity, allowing data recovery if either one or two disks are down in a Storage Pool. ***For maximum data protection, DataDirect Networks recommends the use of RAID 6***.

- Chunk Size
  The chunk size (in KiB blocks) defines the amount of data written to a single disk before proceeding to the next disk in the Storage Pool.

> **NOTE :** RAID 1 is a two member RAID set where the data is mirrored on each disk. There is no parity, hence, the chunk size is fixed.

- Disk Count
  A RAID 1 storage pool consists of 2 physical disks. A RAID 5 Storage Pool may consist of 5 or 9 disks. A RAID 6 Storage Pool may consist of 6 or 10 disks. For maximum performance, select disks with the same characteristics (such as SAS/SATA, capacity, and RPM).

- Drive Type
  Drive types can be SAS or SATA.

- Drive Size
  Drive size is the capacity of the disk drive in GB.

- Spindle Speed
  You may choose (RPM): 15000, 10000, 7200, 5400, or SSD.

- SATAssure
  SATAssure technology is designed by DDN to improve the reliability of enterprise SATA disks and make sure that data integrity is always mentioned for all IO operations. A license is required to run this feature.

To create a storage pool, enter the command:
CREATE POOL RAID_LEVEL=[1|5|6] CHUNK_SIZE=[32|64|128|256] {ASSIGN_POLICY=[SAS|SATA] NUMBER=[2|5|6|9|10] or PHYSICAL_DISK=<list of disks>} INIT_ALLOW=<initialization policy> SATASSURE=[NONE|DATA_INTEGRITY_FIELD|PARITY]

where <initialization policy>=Allow_IO, NO_IO, Priority (1-99, where 99 puts all system resources on the initialization).

You may either explicitly select the disks for the pool by using "**PHYSICAL_DISK=**" or specify the number of disks in the pool by using "**NUMBER=**" in which case the next available disks will be selected. If you specify the **ASSIGN_POLICY** and **NUMBER** parameters together, only the selected type of disks will be used for the storage pool and you will not need to enter the individual disk name. You may disable the SATAssure feature or enable it to use either the data integrity field option or the parity check on read option.

Examples:

- To use only the specified disks:
  CREATE POOL RAID_LEVEL=RAID5 CHUNK_SIZE=64KB PHYSICAL_DISK=62 63 64 65 66
- To use the next available disks:
  CREATE POOL RAID_LEVEL=RAID6 CHUNK_SIZE=128KB NUMBER=6
- To use only the selected type of disks:
  CREATE POOL RAID_LEVEL=RAID6 CHUNK_SIZE=128KB NUMBER=6 ASSIGN_POLICY=SAS

3.3.3 Storage Pool Initialization

When a storage pool is created, initialization begins automatically and will continue in the background until it is completed.

You may use the **SHOW POOL** command to check if the storage pool is being initialized (see Section 3.3.1, "Display Storage Pool Information" for more information).

You may monitor the initialization progress using the **SHOW JOB** command (see Section 3.10.4, "Background Job Priority" for more information).

3.3.3.1 Initialization Job Failure

Rarely, a job may fail to initialize. If so, the event will be reported as:

```
000737 2010-01-25 20:42:09:0684480  G=61 S=2 T=1 RP=0 VP=1
LOG_ST_SET_FAILED  STATE VIRTUAL DISK SET FAILED CONDITION INDEX:00000003 POOL:00000000 INIT    FAILED
```

If an initialization job fails, enter the command **VERIFY POOL <pool-id> FORCE** to restart the initialization.

Examples:

- **SHOW POOL *** command under failed initialization conditions:

```
                                          |Total |Free  |Max   |                 |Disk|  Global  |Spare |
Idx|Name       |State     |Chunk|Raid| Faults |cap GB|cap GB|VD GB | Settings  | Jobs |T/O |spare pool|Policy|
-------------------------------------------------------------------------------------------------------------
  0 pool-0      FAILED      128   6        I      720      0      0 DWMRFI              1  UNASSIGNED  AUTO
  1 pool-1      NORMAL      128   6               720      0      0 DWMRFI              1  UNASSIGNED  AUTO
```

- **SHOW POOL 0 ALL** command under failed initialization conditions:

```
Index:                0
OID:                  0x1a520000
Type:                 STORAGE
Name:                 pool-0
Chunk Size:           128KB (0x100 blocks)
Block Size:           512
RAID Type:            RAID6
Free Raid6 Capacity:  0 GB
Max VD Size:          0 GB
Total Capacity:       720 GB
UUID:                 60001ff0800a3000000000001a520000
Global Spare Pool:    UNASSIGNED
DiskTimeout(FRT):     1 minute
Init Policy:          ALLOW_IO
  .
  .
  .
Initializing:         FAILED
Rebuilding:           FALSE
Paused:               FALSE
  .
  .
  .
```

- **SHOW VD *** command under failed initialization conditions:

```
                                                                |      Home       |Background|
Idx|Name        | State  |Pool|Raid|Cap GB|Settings| Jobs    |Current|Preferred|    Job    |
-------------------------------------------------------------------------------------------
  0 vd-0_0       FAILED     0   5        8    W M I            1(L) 0   1(L) 0   INACTIVE
  1 vd-1_0       FAILED     0   5        8    W M I            1(L) 0   1(L) 0   INACTIVE
  2 vd-2_0       FAILED     0   5        8    W M I            1(L) 0   1(L) 0   INACTIVE
  3 vd-3_1       READY      1   5        8    W M I            1(L) 0   1(L) 0   INACTIVE
```

- **SHOW VD 3 ALL** command under failed initialization conditions:

```
Index:            0
OID:              0x8a680000
Name:             vd-0_0
Pool Index:       0
Pool OID:         0x1a520000
Capacity:         472 GB
Offset:           0x0
State:            FAILED INITIALIZATION
Raidlevel:        RAID6
IO ROUTING:       TRUE
WBC:              TRUE
MIRRORED:         TRUE
Initializing:     FAILED
Paused:           FALSE
AutoWriteLock:    FALSE
```

3.3.4 SATAssure

You may enable SATAssure for a storage pool.

To display the current setting, enter the command (Figure 77):

SHOW POOL=<pool-id> ALL_ATTRIBUTES

Figure 77. Show Storage Pool Details Example Screen

```
***************
*   Pool(s)   *
***************
  Index:                0
  OID:                  0x1a520000
  Type:                 STORAGE
  Name:                 pool-0
   .
   .
   .
  Assign Policy:
    Device Type:        SAS
    Rotation Speed:     NA
    Raw Capacity:       136 GB
  SATAssure:            NONE
   .
   .
   .
```

To enable SATAssure, enter the command:

SET POOL=<pool-id> SATASSURE=[PARITY|DATA_INTEGRITY_FIELD]

where **PARITY** is the parity check on read option and **DATA_INTEGRITY_FIELD** is the integrity field option.

To disable SATAssure, enter the command:

SET POOL=<pool-id> SATASSURE=NONE

> **NOTE :** When the SATAssure setting is changed from **PARITY** to a different setting, an initialization job is started automatically and runs in the background.

3.3.5 Verifying a Storage Pool

Pool verification is a feature that goes with SATAssure. It allows you to run a background job that finds and fixes parity issues (when using the **FORCE** option). It is also automatically initiated when SATAssure encounters an error in an attempt to fix the errors.

To set this mode, use the command:

SET SUBSYSTEM VERIFY_POLICY=<0..99>

where the number range indicates the internal resources to use for the process. Zero is used to turn the feature OFF, any other number will turn the feature on.

To use the "one-time" verify function, use the command:

VERIFY POOL [<pool-id>|*] FORCE_CONSISTENCY

3.3.6 Naming a Storage Pool

The **SET POOL=<pool-id> NAME="<name>"** command lets you specify a name to identify the storage pool (Figure 78).

Figure 78. Set Pool Name Example Screen

```
RAID[0]$ set pool 0 name="RAID 5 Set"
POOL 0 OID: 0x19b40000 attributes set with STATUS='Success' (0x0)
```

3.3.7 Deleting a Storage Pool

The **DELETE POOL=<id>** command deletes the specified storage pool from the system.

The DELETE POOL command erases all the data on the storage pool. You cannot delete a storage pool if there are virtual disks still present on the pool.

3.3.8 Locate a Storage Pool

The **LOCATE POOL=<id>** command provides a visual indication of the specified storage pool. The status LED of the disk modules in the specified storage pool will blink.

3.4 Virtual Disk Management

A virtual disk is the storage unit presented to any attached host. A virtual disk can be created to use all or just a part of the capacity of a single storage pool. Virtual disks allocate space in 8 GiB increments. For example, 16 GiB of storage space will be allocated when creating a virtual disk of 10 GiB.

3.4.1 Display Virtual Disk Information

The **SHOW VIRTUAL_DISK *** command displays the list of configured virtual disks in the system (Figure 79).

Figure 79. Show Configured Virtual Disks Example Screen

```
RAID[0]$ show vd *

                                                                                |          Home    |Background|
Idx|Name            | State  |Pool|Raid|Cap GB|Settings| Jobs  |Current|Preferred|    Job    |
-------------------------------------------------------------------------------------------------------
  0 vd-0_0          READY     0   5     3632   W   I          0(L) 0   0(L) 0   INACTIVE
  1 vd-1_1          READY     1   5     3632   W   I          0(L) 0   1(R) 0   INACTIVE
  2 vd-2_2          READY     2   5     3632   W   I          0(L) 0   0(L) 0   INACTIVE
  3 vd-3_3          READY     3   5     5480   W   I          0(L) 0   1(R) 0   INACTIVE
  4 vd-4_4          READY     4   5     5480   W   I          0(L) 0   0(L) 0   INACTIVE
  6 vd-6_5          READY     5   5     5480   W   I          0(L) 0   1(R) 0   INACTIVE

Total Virtual Disks: 6
```

The **SHOW VIRTUAL_DISK=<id> ALL_ATTRIBUTES** command displays the detailed information of the specified virtual disk (Figure 80).

Figure 80. Show Virtual Disk Attributes Example Screen

```
RAID[0]$ show vd 0 all

  Index:              0
  OID:                0x8a680000
  Name:               vd-0_0
  Pool Index:         0
  Pool OID:           0x1a520000
  Capacity:           472 GB
  Offset:             0x0
  State:              READY
  Raidlevel:          RAID6
  IO ROUTING:         TRUE
  WBC:                TRUE
  MIRRORED:           TRUE
  Initializing:       FALSE
  Paused:             FALSE
  AutoWriteLock:      FALSE
  Critical:           FALSE
  Forced Write-thru:FALSE
  Current Home:       0x38000000  0x00000000 (LOCAL)
  Future Home:        0xffffffff  0x00000000
  Preferred Home:     0x38000000  0x00000000 (LOCAL)
  Job OID:            INACTIVE
  UUID:               60001ff0800a3000000000008a680000

 Total Virtual Disks: 1
```

3.4.2 Creating a Virtual Disk

To create a virtual disk, enter the command:
`CREATE VIRTUAL_DISK CAPACITY=<capacity> POOL=<pool-id>`
where <capacity> is the capacity of the virtual disk in GiB or type "max" to use all available capacity; <pool-id> is the storage pool to be used by this virtual disk (Figure 81).

A message is displayed to indicate whether the virtual disk creation was successful.

Figure 81. Create Virtual Disks Example Screen

```
RAID[0]$ create virtual_disk capacity=16 pool=0
VIRTUAL_DISK 0 OID=0x89ba000 creation STATUS='Success' (0x0)

RAID[0]$ create virtual_disk capacity=32 pool=0
VIRTUAL_DISK 1 OID=0x89bb001 creation STATUS='Success' (0x0)
```

3.4.3 Naming a Virtual Disk

The **`SET VIRTUAL_DISK=<id> NAME=<name>`** command lets you specify a name to identify the virtual disk (Figure 82).

Figure 82. Set Virtual Disk Name Example Screen

```
RAID[0]$ set vd 0 name=LUN ”

VIRTUAL_DISK 0 OID=0x89ba000 attributes set with STATUS=’Success’ (0x0)
```

3.4.4 Deleting a Virtual Disk

The **`DELETE VIRTUAL_DISK=<id>`** command deletes the specified virtual disk from the system. However, it does not delete presentations for that virtual disks. If a presentation is configured for a virtual disk, deleting the virtual disk will return an error. You must first delete all the presentations configured on that virtual disk (see Section 3.6.3, "Presentation Commands").

When you delete a virtual disk, you lose access to all data that was stored on that virtual disk.

3.4.5 Presentation to Hosts (not applicable on SFA10000E)

Virtual disks are only presented to the hosts that have been given authorized access. Refer to Section 2.7.12, "Present Virtual Disk to External Host (SFA10000 Only)" for detailed instructions on configuring virtual disk presentations to hosts. Refer to Section 3.6, "Presentations" for more information on commands related to presentation.

3.5 Spare Pool Management

The SFA10000 supports the concept of spare pool which contains Physical Disks that can be used as spare disks. spare pools can be shared by storage pools or dedicated to a particular storage pool in which case the spare disks will only be used by the designated storage pool. After a spare pool is created, you will need to select disk(s) from the Unassigned Pool and assign it the spare pool.

In the event of disk failure, a disk that is assigned to a spare pool is automatically swapped into the storage pool to replace the failed disk. A rebuild then occurs immediately, reducing the possibility of data loss due to additional disk failures.

NOTE : Each storage pool should have a spare pool assigned to it. If a storage pool's attributes show "Global Spare Pool: UNASSIGNED", no spare pool will serve this storage pool.

3.5.1 Display Spare Pool Information

Use the **SHOW SPARE_POOL *** command to display the list of configured spare pool(s) (Figure 83).

Figure 83. Show Configured Spare Pools Example Screen

```
RAID[0]$ show spare_pool *

                                  |Disk |Total |Total|Storage |
Idx|Name                  |Blocks| T/O |Cap GB| PDs |Pool Idx|
-------------------------------------------------------------
  6 spare_pool-6            512    10     528    2
  7 spare_pool-7            512    10     528    2

Total Spare Pools: 2
```

To display the detailed information of the spare pool, enter the command:
SHOW SPARE_POOL=<id> ALL_ATTRIBUTES command (Figure 84).

Figure 84. Show Spare Pool Attributes Example Screen

```
RAID[0]$ show spare_pool 6 all

  Index:               6
  OID:                 0x19bc0006
  Type:                GLOBAL SPARE
  Name:                19bc0006
  Block Size:          0x200
  DiskTimeout(FRT):    10 minutes
  Total Capacity:      1409024 MBs
  UUID:                0x00
  Total Phy Disks      2
  Storage Pool List:   Index: 0  OID:0x1c790000
```

To display the list of disks in the spare pool, enter the command (Figure 85):
SHOW SPARE_POOL=<id> PHYSICAL_DISKS

Figure 85. Show Spare Pool Physical Disks Example Screen

```
RAID[0]$ show spare_pool 6 pd
                                                                                |Health|
Encl|Slot| Vendor |     Product ID     |Type|Cap GB | RPM|Revision|    Serial Number   |Pool|State | Idx|State |       WWN       |
-------------------------------------------------------------------------------------------------------------------------------------
   2   53 HITACHI  HUS153030VLS300       SAS    279  15K A410     JHVYKUKC5303          14  SPARE  648 NORM    5000cca00d35c6b0
   2   58 HITACHI  HUS153030VLS300       SAS    279  15K A410     JHVWU54C5303          14  SPARE  652 NORM    5000cca00d3291f4

Total Physical Disks:      2
Total Assigned Disks:      2
Total Unassigned Disks:    0
  Total SAS Disks:         2
  Total Member State Normal:   2
```

3.5.2 Creating a Spare Pool

Use these commands to create a spare pool, add disk to the spare pool, and assign the spare pool to a storage pool:

- **CREATE SPARE_POOL**
- **ASSIGN PHYSICAL_DISK=<disk-id> TO_POOL <spare-pool-id>**
- **SET POOL=<pool-id> SPARE_POOL=<spare-pool-id>**
  where <pool-id> and <spare-pool-id> are the index names of the storage pool and spare pool respectively.

3.5.3 Naming a Spare Pool

The **SET SPARE_POOL=<pool-id> NAME="<name>"** command lets you specify a name to identify the spare pool.

3.5.4 Deleting a Spare Pool

To delete a disk from the spare pool, enter the command:

ASSIGN PHYSICAL_DISK=<disk-id> TO_POOL 0x1800ffff

To delete a spare pool, enter the command:

DELETE SPARE_POOL=<pool-id>

3.6 Presentations

NOTE : For SFA10000E, this feature is not applicable.

A Presentation on a SFA10000 has the following components:

- Discovered Initiators
  A Discovered Initiator is defined as a Fibre Channel or InfiniBand initiator that has logged into the SFA10000. This is normally a single port on a Fibre Channel or InfiniBand HCA in an external computer.

- Host
  A Host is an abstract client of a virtual disk that has all the access to all data within the virtual disk. This most commonly represents an external computer with one or more Fibre Channel or InfiniBand initiators. Host attributes are OS (operating system) of which you can select GENERIC, WINDOWS, LINUX or MAC_OS. The OS attribute dictates the way the controller responds to certain Fibre Channel commands where appropriate. For example, Mac OS X expects a different response from an inquiry of LUN0 (the controller LUN) than a Windows host.

  - OSTYPE: “Standard” uses current settings and “Custom” allows for requests for different adjustments from OEMs.

- Channel
  A Channel is one of the Fibre Channel or InfiniBand ports on the controller. Channel attributes are MODE of which you can select either MAC_OS or STANDARD. When MAC_OS is selected, the port’s node name will be set differently in order to be visible to a Macintosh system.

- Stack
  A Stack is the input side of the subsystem. In the case of the SFA10000, the stack is defined as the Fibre Channel interface. In future versions of the product, different interface stacks may exist.

- Presentation
  A Presentation is the relationship between a Host and a virtual disk. A Presentation implies that the related Host has some sort of access to the virtual disk. Attributes of a Presentation are:
  - PORT: from which the host will see the virtual disk
  - READ_ONLY: controls read only access
  - PRESENT_HOME_ONLY: presents the specified virtual disk from its designated home controller only
  - LUN: user-specified LUN number that the virtual disk will show to the host.

3.6.1 Discovered Initiator Commands

To map a host to a discovered initiator, use the command:
APPLICATION IMPORT DISCOVERED_INITIATOR=<initiator-id> HOST=<host-id>

To display the currently available initiators, use the command:
APPLICATION SHOW DISCOVERED_INITIATORS *

To display the attributes of a specified initiator, use the command:
APPLICATION SHOW DISCOVERED_INITIATORS=<initiator-id> [ALL ATTRIBUTES]

To create an application initiator for the specified host, use the command:
APPLICATION CREATE INITIATOR HOST=<host-id> WWPN=<integer>

To delete the specified initiator, use the command:
APPLICATION DELETE INITIATOR=<initiator-id>

3.6.2 Host Commands

To create a host, use the command:
APPLICATION CREATE HOST [OSTYPE=GENERIC|LINUX|MAC_OS|WINDOWS]

To delete the specified host, use the command:
APPLICATION DELETE HOST=<host-id>

To display the presentation associated with the specified application host, use the command:
APPLICATION SHOW HOST=<host-id> [PRESENTATIONS]

To display the initiators associated with the specified application host, use the command:
APPLICATION SHOW HOST=<host-id> [INITIATORS]

3.6.3 Presentation Commands

To present a virtual disk to the specified host, use the command:
APPLICATION CREATE PRESENTATION HOST=<host-id> VIRTUAL_DISK=<vd-id> LUN=<LUN-id>

To delete a presentation, use the command:
APPLICATION DELETE PRESENTATION HOST=<host-id> VIRTUAL_DISK=<vd-id> [FORCE]
Note the optional parameter of **FORCE** deletes without confirmation. The wildcard ***** deletes all presentations.

To set the specified attribute to the specified value, use the command:
APPLICATION SET PRESENTATION=<object-id> [<attribute-name>=<value>]
Attributes are:

ENABLE=ALL|NONE|<channel-id> where the channel-id is the object -id of an Enabled Client Channel.

LUN=<integer> where the integer is a Logical Unit Number (LUN) that will be used to present the associated virtual disk to the associated Host. NOTE: Each LUN integer entered for a presentation is cumulative and does not replace the previous entry.

HOME_ONLY=[TRUE|FALSE] which enables/disables the home_only parameter.

READ_ONLY=[TRUE|FALSE] which enables/disables the read_only parameter.

3.6.3.1 Persistent Reservation Support

SCSI-2 Persistent Reservation commands support is included for legacy host file system. SFA OS receives either a SCSI-2 RESERVE/RELEASE or SCSI-3 PRIN/PROUT and responds with the proper type, SCSI-2 or SCSI-3. SFA OS will not allow SCSI-2 and SCSI-3 reservation styles to be mixed on the same LUN. Attempting to do so will return reservation conflicts.

Note that a promiscuous host may not be used with persistent reservations. In the following command, if **<host-id>** is **ALL**, you have a promiscuous host.

```
APPLICATION CREATE PRESENTATION VD=<VD-id> HOST=<host-id>
```

SFA OS V1.4.0.0 includes the following commands:

- **SCSI-3 Persistent Reserve In**:
  PRIN Read Keys (00h) MS
  PRIN Read Reservation (01h) MS
  PRIN Report Capabilities (02h)
  PRIN Read Full Status (03h)
- **SCSI-3 Persistent Reserve Out**:
  PROUT Reserve (01h), Scope=LU_SCOPE(0h), Type=Write Exclusive Registrants Only (5h) MS or Type=Exclusive Access (6h) MS
  PROUT Release (02h) MS
  PROUT Clear (03h) MS
  PROUT Preempt (04h) MS

3.6.4 Additional Configuration Considerations for Mac® Hosts

Apple®/Mac OS® handles Fibre Channels Port/Node WWNs uniquely. Apple hosts expect node names to be unique across all subsystem ports.

All other OS's expect the FC node names to be the same across all subsystem ports.

Because of this uniqueness, VD presentations must be uniquely set up when an Apple/Mac presentation is being made from a SFA10000 subsystem.

> **NOTE :** You MUST create the host object with **OSTYPE=MAC** for Mac hosts.

You must identify at least one controller port (channel) (preferably one controller port for each SFA10000 controller) on the storage subsystem as operating in the Mac mode:

```
APPLICATION SET CHANNEL [0|1|2|3] MODE MAC
```

> **NOTE :** Setting/changing the mode of a "channel" will require a shutdown/restart of the controller for the change to actually take place.

If a Mac presentation is being made from the subsystem, NO other VD from this subsystem should be presented to "ALL" hosts !

In a SAN and/or MIXED OS environment of Mac and other operating systems, create the presentations as follows:

```
APPLICATION CREATE PRESENTATION VD=<vd-id> HOST=<host_index>
ENABLE=NONE
```

This will disable channels through which to make a presentation. Then you will next enable the specific channels through which you wish to make the presentation:

```
APPLICATION SET PRESENTATION <index> ENABLE {0|1|2|3}
```

To enable two channels, it is necessary to enter two separate commands.

In the current implementation, you can disable all channels, enable all channels, or enable one channel at a time through the CLUI. If you wish to have two channels enabled, use the following commands:

```
APPLICATION SET PRESENTATION <index> ENABLE 0
APPLICATION SET PRESENTATION <index> ENABLE 2
```

Channel 0 is C0RP0. Channel 1 is C0RP1. These are the two RAID processors on controller 0.

Channel 2 is C1RP0. Channel 3 is C1RP1. These are the two RAID processors on controller 1.

If you have an existing storage configuration running Windows/Linux presentations and you wish to ADD a VD presentation for MAC:

- You will have to sustain a storage outage now so that you can disable “all” VD presentations thru any channel; and
- Re-enable the appropriate channels one at a time that will be serving VDs to the Windows/Linux hosts.
- You will then set up the appropriate channel(s) that will operate in MAC mode.
- At this point you can restart the storage subsystem, proceed in creating your MAC VD presentations and enable them through their separate channels.

3.6.5 Configure Presentations of Virtual Disks to Hosts

NOTE : If you are configuring a storage subsystem with any “MAC OSx presentations”, it is absolutely required that you read and understand Section 3.6.4, "Additional Configuration Considerations for Mac® Hosts" prior to actually doing your configuration.

There are three objects/relationships that must be properly established in order to create a presentation of a LUN (virtual disk) to a host:

- **Establish a host object** for EACH host that wishes access to the SFA10000 storage system virtual disk.

NOTE : On the storage subsystem you ONLY need to create a single host object for any/all virtual disk presentations to that host.

- **Import a discovered initiator** into an established relationship with a host object. An association between a discovered initiator FC Port/Node WWN (World Wide Name) to a host object is established in this operation. This association of host/port WWN to host object is maintained persistently within the subsystems configuration information along with all other information. If the host to controller association is subsequently disrupted and then re-established, the controller is able to maintain this relationship until such

time that the configuration in the storage controller is cleared or that relationship is deleted.

TIP : Use of a host based HBA utility such as HBAnyware® or SANsurfer will allow you to examine the FC Port/Node WWNs on the host in the easiest manner.

- **Assign a virtual disk to a host object.**

NOTE : You may ONLY present a virtual disk ONCE to the same host object.

You can present a virtual disk to multiple hosts; however, this is dangerous. If doing so, presenting them as Read-Only to the other hosts would be appropriate.

NOTE : You may present a virtual disk to ALL hosts. This may be appropriate in some limited system environments; however, it is best practice to not MIX the presentations, where some VDs are selectively presented while others are presented all. Managing at the host end may become confusing.

To present a virtual disk to ALL hosts (promiscuous mode), use the command:
`APPLICATION CREATE PRESENTATION VD=<VD-id> HOST=ALL`

NOTE : The promiscuous mode may not be used with persistent reservations. The above setting will cause SCSI2/3 Reservation commands to fail.

3.6.5.1 Create Host Object

The following example demonstrates presenting 6 virtual disks to 3 separate hosts (2 Windows hosts and a Linux host) from a single storage subsystem. Although this configuration will have a FC Switch, the steps are identical.

NOTE : In the example below, the CLUI is operating in the RAID subject mode. The ASM (Application Stack Management) commands must be preceded by the subject application.

Examples:

```
RAID[1]$ application create host name=co-ls1 ostype=Linux
HOST 301 OID=0x1d5e012d creation STATUS='Success' (0x0)

RAID[1]$ application create host name=co-test-d10 ostype=windows
HOST 302 OID=0x1d5f012e creation STATUS='Success' (0x0)

RAID[1]$ app create host name=co-test-d08 ostype=windows
HOST 303 OID=0x1d60012f creation STATUS='Success' (0x0)
```

```
RAID[1]$ application show host *

                                | Stack |   Host Mode Attributes    |
Index | Host Name               | Index | OS Type | Characteristics |
-------------------------------------------------------------------------
00304   co-ls1                  00000    LINUX     0x0000000000000001
00305   co-test-d10             00000    WINDOWS   0x0000000000000001
00306   co-test-d08             00000    WINDOWS   0x0000000000000001

Total Hosts: 3
```

Recommendations:

- When creating host objects, always include a descriptive “name” for the host objects that is easily recognizable. This benefits the storage administrator in managing storage connectivity issues.
- The **OS_TYPE** determines the FC flow control mechanism for host I/O.
- The Linux OS type is currently appropriate for all versions/implementations of Linux and UNIX.
- The Windows OS type is currently appropriate for all versions of the Windows operating systems.
- It is imperative the **OS_TYPE** for Apple/MAC hosts be appropriately set to MAC. The setup requirements involving MAC OSX hosts require additional steps and considerations. Please see Section 3.6.4, "Additional Configuration Considerations for Mac® Hosts".

3.6.5.2 Identifying Host FC Connections via Ports

The host ports of the controller dynamically acknowledge and log in any FC Host/Port WWN connection that it can sync up with. When you query the storage subsystem for the discovered initiators, the listing you get will be the current set of connections that are logged in. If a connection is broken and then re-established, the re-discovered Port/Node WWN will be assigned a new host_index number. This is inconsequential after a specific Port/Node WWN is directly associated with a specific host that is done in a subsequent step.

Example of discovered initiators:

```
RAID$ app show discovered_initiator *

                  |          Initiator Identifier          | Controller 0  | Controller 1  |
Index |Type|  ID  |    node            |      port         | RP 0  | RP 1  | RP 0  | RP 1  |
-------------------------------------------------------------------------------------------
00003   FC |0x010600 0x20000000c9813cc9 0x20000000c9813cc9          |       |     3 |
00005   FC |0x010800 0x20000000c9813cc8 0x20000000c9813cc8          |       |       |0
00006   FC |0x010900 0x20000000c9813a47 0x20000000c9813a47          |       |    2  |
00007   FC |0x010a00 0x20000000c9813a46 0x20000000c9813a46          |       |       |  1
00011   FC |0x010300 0x2000001b32827e95 0x2000001b32827e95          |0      |       |
00012   FC |0x010200 0x2000001b32a27e95 0x2000001b32a27e95          |0      |       |

Total FC Initiators: 6
```

In order to provide an association between a host and a virtual disk, you must identify the FC connection (its Node/Port WWN name to the host object). You may use the following techniques:

- Utilize the physical sticker tag information on the HBA.

- Connect a single host at a time (and all its FC connections to the subsystem) to the subsystem (direct connect or via a switch); identify their WWNs as a discovered initiator; then a proceed to plug in subsequent hosts, noting their WWNs.
- Utilize a host based FC Adapter utility such as HBAnyware (for Emulex HBAs) (Figure 86) or SANsurfer (for QLogic HBAs) (Figure 87).

Figure 86. HBAnyware Screen

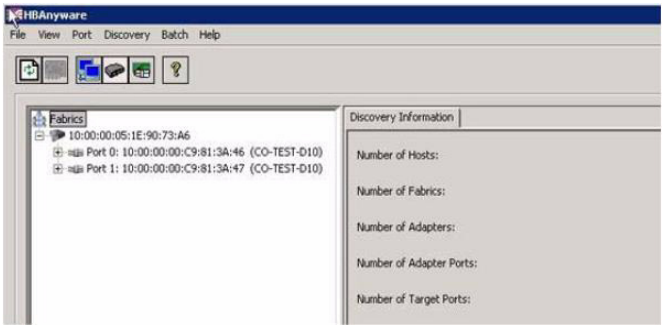

Figure 87. SANsurfer FC HBA Manager Screen

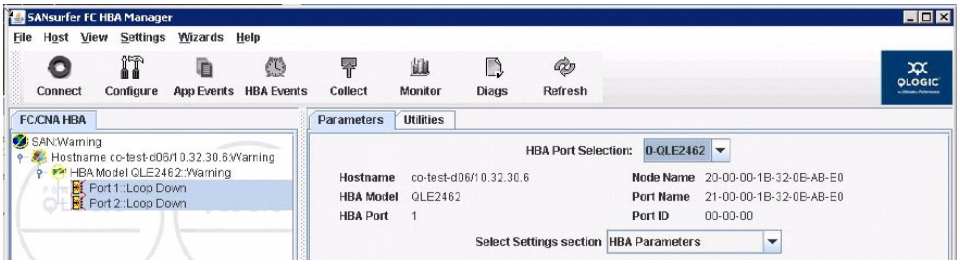

Example for Linux system with driver loaded (example of QLogic):

```
# cd /proc/scsi/qla2xxx
# ls
1 2
# grep adapter-port 1
scsi-qla0-adapter-port=210000e08b9d6149;

bm-dell-09:/proc/scsi/qla2xxx # grep adapter-port 2

scsi-qla1-adapter-port=210100e08bbd6149;

# lsscsi -g | grep -i ddn
[1:0:0:1]   disk   DDN    S2A 8500      5.31 /dev/sdc /dev/sg2
[2:0:0:1]   disk   DDN    S2A 8500     5.31 /dev/sdg /dev/sg6
```

Example for Linux system without driver loaded:

```
# cd /sys/class/fc_host
# ls -la
total 0
drwxr-xr-x 4 root root 0 Jul 15 11:47 .
drwxr-xr-x 43 root root 0 Jul 15 11:48 ..
drwxr-xr-x 3 root root 0 Jul 15 11:48 host3
drwxr-xr-x 3 root root 0 Jul 15 11:48 host4
# cd host3
# ls -la
total 0
drwxr-xr-x 3 root root 0 Jul 15 11:48 .
drwxr-xr-x 4 root root 0 Jul 15 11:47 ..
lrwxrwxrwx 1 root root 0 Jul 15 11:47 device ->
../../../devices/pci0000:00/0000:00:04.0/0000:0a:00.0/host3
-r--r--r-- 1 root root 4096 Jul 17 15:31 fabric_name
--w------- 1 root root 4096 Jul 17 15:31 issue_lip
-r--r--r-- 1 root root 4096 Jul 17 15:31 node_name
-r--r--r-- 1 root root 4096 Jul 17 15:31 port_id
-r--r--r-- 1 root root 4096 Jul 17 15:31 port_name
-r--r--r-- 1 root root 4096 Jul 17 15:31 port_state
-r--r--r-- 1 root root 4096 Jul 17 15:31 port_type
-r--r--r-- 1 root root 4096 Jul 17 15:31 speed
drwxr-xr-x 2 root root 0 Jul 15 11:47 statistics
lrwxrwxrwx 1 root root 0 Jul 15 11:48 subsystem -> ../../../class/fc_host
-r--r--r-- 1 root root 4096 Jul 17 15:31 supported_classes
-r--r--r-- 1 root root 4096 Jul 17 15:31 symbolic_name
-rw-r--r-- 1 root root 4096 Jul 17 15:31 system_hostname
-rw-r--r-- 1 root root 4096 Jul 17 15:31 tgtid_bind_type
--w------- 1 root root 4096 Jul 15 11:47 uevent
# cat port_name
0x2100001b3282dc50
```

NOTE : In switch environments you may disable/enable the ports on the switch to identify which physical connection you are dealing with.

Utilizing one or more of the techniques above, you can document an association between the physical host and the discovered initiators on the storage.

Example of showing imported initiators:

```
RAID[1]$ application show initiator *

             | Host    |          Initiator Identifier       |
Index | Type | Index   |   node                |     port          |
-------------------------------------------------------------------
00004   FC    0x010600  0x20000000c98107cb 0x10000000c98107cb <-- co-test-d08
00006   FC    0x010900  0x20000000c9813a47 0x10000000c9813a47 <-- co-test-d10
00007   FC    0x010a00  0x20000000c9813a46 0x10000000c9813a46 <-- co-test-d10
00008   FC    0x010300  0x20000000c98107ca 0x10000000c98107ca <-- co-test-d08
00011   FC    0x010200  0x2000001b32827e95 0x2100001b32827e95 <-- co-ls1
00012   FC    0x010800  0x2001001b32a27e95 0x2101001b32a27e95 <-- co-ls1
```

3.6.5.3 Import Discovered Initiators

Now identify the WWN and the discovered initiator index number with the host and its index number and import the appropriate discovered initiator to an association with the appropriate host.

Example of importing discovered initiators:

```
RAID[1]$ application import discovered 11 host 304
INITIATOR 12 OID=0x2d66000c imported from DISCOVERED_INITIATOR 11 OID=0x3000000b
STATUS='Success' (0x0)

RAID[1]$ application import discovered 12 host 304
INITIATOR 13 OID=0x2d67000d imported from DISCOVERED_INITIATOR 12 OID=0x3000000c
STATUS='Success' (0x0)

RAID[1]$ application import discovered 4 host 306
INITIATOR 14 OID=0x2d68000e imported from DISCOVERED_INITIATOR 4 OID=0x30000004
STATUS='Success' (0x0)

RAID[1]$ application import discovered 8 host 306
INITIATOR 15 OID=0x2d69000f imported from DISCOVERED_INITIATOR 8 OID=0x30000008
STATUS='Success' (0x0)

RAID[1]$ app import discovered=6 host=305
INITIATOR 16 OID=0x2d6a0010 imported from DISCOVERED_INITIATOR 6 OID=0x30000006
STATUS='Success' (0x0)

RAID[1]$ app imp disc 7 host 305
INITIATOR 17 OID=0x2d6b0011 imported from DISCOVERED_INITIATOR 7 OID=0x30000007
STATUS='Success' (0x0)

RAID[1]$ application show initiator *

             | Host  |            Initiator Identifier       |
Index | Type | Index |      node name     |   port name      |
-------------------------------------------------------------
00012   FC    00304    0x2000001b32827e95  0x2100001b32827e95
00013   FC    00304    0x2001001b32a27e95  0x2101001b32a27e95
00014   FC    00306    0x20000000c98107cb  0x10000000c98107cb
00015   FC    00306    0x20000000c98107ca  0x10000000c98107ca
00016   FC    00305    0x20000000c9813a47  0x10000000c9813a47
00017   FC    00305    0x20000000c9813a46  0x10000000c9813a46

Total FC Initiators: 6
```

3.6.5.4 Present a Virtual Disk to a Host

From your list of virtual disks that have been created, identify the host that you wish to present this LUN to and create the presentation.

Recommendations:

- Commands entered through the CLI may take a shortened form of the unique word as demonstrated in the above examples
- Spaces may be used in place of "=" sign in command syntax

Example of list of candidate virtual disks for presentation:

```
RAID[1]$ show vd *
                                                                    |      Home       |Background|
Idx|  Name    | State |Pool|Raid|Cap GB|Settings| Jobs |Current|Preferred|   Job    |
-------------------------------------------------------------------------------------------------
 49 vd-49_29   READY   29   6   1400     W I          1(L) 0   1(L) 0   INACTIVE
 50 vd-50_30   READY   30   6   1400     W I          0(R) 0   0(R) 0   INACTIVE
 51 vd-51_31   READY   31   1    344     W I          1(L) 0   1(L) 0   INACTIVE
 52 vd-52_32   READY   32   1    344     W I          0(R) 0   0(R) 0   INACTIVE
 53 vd-53_33   READY   33   5   1400     W I          1(L) 0   1(L) 0   INACTIVE
 54 vd-54_34   READY   34   5   1400     W I          0(R) 0   0(R) 0   INACTIVE
 55 vd-55_35   READY   35   6   3608     W I          1(L) 0   1(L) 0   INACTIVE
 56 vd-56_36   READY   36   6   3608     W I          0(R) 0   0(R) 0   INACTIVE
 57 vd-57_37   READY   37   5   3608     W I          1(L) 0   1(L) 0   INACTIVE
 58 vd-58_38   READY   38   5   3608     W I          0(R) 0   0(R) 0   INACTIVE
 59 vd-59_39   READY   39   1    344     W I          1(L) 0   1(L) 0   INACTIVE
 60 vd-60_41   READY   41   1    896     W I          1(L) 0   1(L) 0   INACTIVE

Total Virtual Disks: 12
```

Example of creating VD presentation to Windows host co-test-d10 (not specifying SCSI LUN_ID):

The next two VDs are presented to a Windows host and the subsystem assigns the SCSI LUN_ID.

```
RAID[1]$ application create presentation vd 51 host 305
PRESENTATION 2482 OID=0x256e09b2 creation STATUS='Success' (0x0)

RAID[1]$ application create presentation vd 52 host 305
PRESENTATION 2483 OID=0x256f09b3 creation STATUS='Success' (0x0)

RAID[1]$ application show presentation *
                                                          |      Channel Mask         |
                                                          | Controller 0|Controller 1 |
Pres. | Host            Host |   VD  |     |Home |Read | RP 0 | RP 1 | RP 0 | RP 1 |
Index | Name           Index | Index | LUN |Only |Only | 0123 | 0123 | 0123 | 0123 |
-----------------------------------------------------------------------------------------
02480  co-ls1          00304   00049   060   OFF   R/W   ffff   ffff   ffff   ffff
02481  co-ls1          00304   00050   061   OFF   R/W   ffff   ffff   ffff   ffff
02482  co-test-d10     00305   00051   000   OFF   R/W   ffff   ffff   ffff   ffff
02483  co-test-d10     00305   00052   001   OFF   R/W   ffff   ffff   ffff   ffff

Total Presentations: 4
```

NOTE : For Windows environments, you should let the subsystem specify the LUN_ID for a virtual disk.

```
RAID[1]$ application create presentation vd 58 host 306 lun 10
PRESENTATION 2484 OID=0x257009b4 creation STATUS='Success' (0x0)

RAID[1]$ application create presentation vd 59 host 306 lun 12
PRESENTATION 2485 OID=0x257109b5 creation STATUS='Success' (0x0)
```

```
RAID[1]$ app show pres *

                                                                |      Channel Mask         |
                                                                | Controller 0|Controller 1 |
Pres. | Host            Host |   VD  |     |Home |Read | RP 0 | RP 1 | RP 0 | RP 1 |
Index | Name           Index | Index | LUN |Only |Only | 0123 | 0123 | 0123 | 0123 |
-------------------------------------------------------------------------------------
02480  co-ls1          00304   00049   060   OFF   R/W   ffff   ffff   ffff   ffff
02481  co-ls1          00304   00050   061   OFF   R/W   ffff   ffff   ffff   ffff
02482  co-test-d10     00305   00051   000   OFF   R/W   ffff   ffff   ffff   ffff
02483  co-test-d10     00305   00052   001   OFF   R/W   ffff   ffff   ffff   ffff
02484  co-test-d08     00306   00058   010   OFF   R/W   ffff   ffff   ffff   ffff
02485  co-test-d08     00306   00059   012   OFF   R/W   ffff   ffff   ffff   ffff

Total Presentations: 6
```

Recommendations:

- At times some hosts may have device discovery issues if there are “gaps” in the SCSI LUN_ID space. Letting the subsystem assign (as it will by default) the lowest SCSI_ID it can from this subsystem, may be considered a best practice.
- Use of a host based HBA utility such as HBAnyware® or SANsurfer greatly enhances the ability to determine perceived connectivity issues between the storage subsystem and the host.
- Your hosts “multipath driver” must be enabled correctly to recognize DDN’s Device/Hardware ID.
  - For Windows Server 2008, the Device/Hardware_ID information must be entered into the MPIO stack as an 8 bit / 16 character space padded field.
  - For Linux OS’s, the /etc/multipath.conf file must be edited, and the Device/Hardware_ID is NOT padded with the space character.

    ```
    Device Hardware_ID
    DDN SFA10000
    ```

3.7 Network Time Protocol Mode

NTP (Network Time Protocol) mode is available on the SFA10000. It provides a means for the controllers to synchronize their time across a network, usually within a small number of milliseconds over a long period of time. You can enter up to four NTP addresses as the time servers.

To enable the NTP mode, enter the command:

```
SET SUBSYSTEM NTP=[<ip address list, up to 4>|NONE]
```

Examples:

```
SET SUBSYSTEM NTP=1.2.3.4 1.2.3.5 3.4.5.6 6.5.7.8
SET SUBSYSTEM NTP=1.2.3.4
SET SUBSYSTEM NTP=NONE
```

Specifying a list of NTP addresses will start the NTP mode on each controller, using that list of NTP addresses as the time servers. Specifying NONE will turn off NTP mode.

To display the current settings, enter the command(Figure 88):

```
SHOW SUBSYSTEM ALL_ATTRIBUTES
```

Figure 88. Show NTP Settings

```
RAID[0]$ show subsystem all

********************
*    Subsystem     *
********************

RP Subsystem Name:      SFA10000
UID:                    60001ff0800a30000000000030000000
Subsystem Time:         Fri Apr 22 06:26:39 2011
Locate Dwell Time:      120 seconds
Enabled Licenses:       RAID6 SATASSURE
Fast Timeout:           ON
Pool Verify Priority:   10%
NTP Mode:               ON
                        (10.32.16.24)
                        (10.32.16.25)
Drive Error Tolerance:  HIGH
Single Controller WB:   ENABLED
Fri Apr 22 06:26:39 2011
```

When in NTP mode:

- The time is set between the two controllers under the following conditions:
  - when the controllers boot and discover each other
  - when a **SET SUBSYSTEM** command is issued, whether it is setting the time or not
- Each controller will attempt to synchronize with the specified NTP servers. If none of the servers are valid, the time on each controller is free to drift independent of any other time source (and independent of the other controller).
- The **SET SUBSYSTEM DATA_AND_TIME** command will fail with a status of “Setting date/time while in NTP mode”.
- Once the Clock code has finished calibration (100 seconds), NTP will be started. Both controllers will have the same NTP settings, and so presumably will have synchronized time once NTP sets the time.
- Issuing a new **SET SUBSYSTEM NTP** command with a new set of IP addresses will stop and restart NTP.

- If NTP finds a time difference of more than 128 msec, it will "jump" the time to the correct time. This will result in a discontinuity in the event log, logdisk, syslog, and anywhere that records a timestamp.
- NTP will always set the controllers to UTC (Coordinated Universal Time). There is no option to set time zones or otherwise change the offset from UTC.

When not in NTP mode:

- The master controller uses its time to set the time on the other controller.
- The **`SET SUBSYSTEM DATE_AND_TIME`** command is used to set the time of both controllers together. The system records time using the military method, which records hours from 00 to 24, not in a.m. and p.m. increments of 1 to 12. Settings are automatically adjusted for leap years.

  To change the system date and time to March 1, 2010 2:15:32 pm, for example, type:

  ```
  SET SUBSYSTEM DATE_AND_TIME=2010:3:1:14:15:32
  ```

- Once the time is set, the time on the two controllers is free to drift independent of the other controller.

3.8 Disk Rebuild

The SFA10000's automatic disk failure recovery procedures ensure that absolute data integrity is maintained while operating in degraded mode.

3.8.1 Full and Partial Rebuilds

Both full and partial (fractional) rebuilds are supported in SFA10000. Rebuilds are done at the storage pool level, not the virtual disk level. Partial rebuilds will reduce the time to return the storage pool to full redundancy and can be controlled by adjustable time limit for each storage pool. Permitted time limits are 0 (in which case partial rebuild is off) to 240 minutes. The timer will begin when a disk in the storage pool is declared missing. If the disk reappears prior to the expiration of the timer, a fractional rebuild will be done. Otherwise, the disk will be declared failed, replaced by a spare and a full rebuild will begin to return the storage pool to full redundancy. The default partial rebuild timer (Disk Timeout) setting is 10 minutes.

Under heavy write workloads, it is possible that the number of stripes that need to be rebuilt will exceed the system's internal limits prior to the timer expiration. When this happens, a full rebuild will be started automatically instead of waiting for the partial rebuild timeout.

Use the **SHOW POOL=<id> ALL_ATTRIBUTES** command to display the current Disk Timeout setting (Figure 89).

Figure 89. Show Pool All Attributes

```
RAID[0]$ show pool 0 all
  Index:                0
  OID:                  0x1a520000
  Type:                 STORAGE
  Name:                 pool-0
  Chunk Size:           128KB (0x100 blocks)
  Block Size:           512
  RAID Type:            RAID6
  Free Raid6 Capacity:  0 GB
  Max VD Size:          0 GB
  Total Capacity:       720 GB
  UUID:                 60001ff0800a3000000000001a520000
  Global Spare Pool:    0x1aca000e  (Index 14)
  DiskTimeout(FRT):     1 minute
    .
    .
    .
```

To change the disk timeout setting, use this command:

SET POOL=<id> DISK_TIMEOUT=<timeout>

where <timeout> is in the range of <0..240> minutes. The default setting is 10.

3.8.2 Sparing Policy

Each storage pool has a sparing policy that determines what happens when a physical disk within the pool fails (or becomes inaccessible). In the event of a disk failure, the SFA10000 will automatically initiate a disk rebuild if the sparing policy is set to automatic and a spare pool has been assigned to the storage pool.

Use the **SHOW POOL=<id> ALL_ATTRIBUTES** command to display the current settings (Figure 90).

Figure 90. Show Pool All Attributes

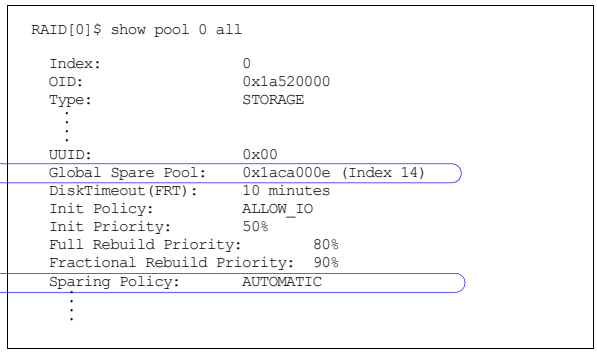

To change the sparing policy setting, use this command:

SET POOL=<id> SPARING_POLICY=[AUTOMATIC|MANUAL]

The default setting is automatic which is recommended.

A rebuild operation can take up to several hours to complete, depending on the size of the disk and rate of rebuild. Refer to Section 3.10, "Performance Management" on page 83 for information on how to adjust the rate of rebuild.

3.8.3 Manual Disk Replace/Rebuild

If a storage pool does not have a spare pool assigned to it, it becomes necessary to manually add a disk to the storage pool to replace a failed disk.

To add a disk to a storage pool to replace a failed disk, use this command:

ASSIGN PHYSICAL_DISK=<id of disk to be added> TO_POOL=<pool-id> SET_SPARE

The disk will be put into the storage pool as the spare disk and it will be used in the storage pool to replace the failed disk.

3.8.4 Manual Fail/Rebuild of a Disk

The **SET PHYSICAL_DISK <disk-id> FAILED** command instructs the system to fail the specified disk. When a non-SPARE disk is specified and it is failing, the disk will not cause a multi-channel failure. The disk is marked as failed. An attempt is made to replace it with a spare disk. When a SPARE disk is specified, it is released, but marked as unhealthy and unavailable.

CLEAR PHYSICAL_DISK <disk-id> FAILED changes the specified disk's "Failed" state to "FALSE". After clearing a disk, use the **ASSIGN PHYSICAL_DISK <disk-id> TO_POOL <pool-id> SET_SPARE** command to add the disk back to the storage pool it was failed from. This command will also initiate a rebuild if a spare has not already been assigned to the storage pool.

3.9 Restarting the SFA10000

3.9.1 System Restart

The **SHUTDOWN CONTROLLER [LOCAL|REMOTE|0|1] RESTART** command performs a restart on the specified controller.

The **SHUTDOWN SUBSYSTEM RESTART** command performs a restart on both controllers.

These commands will prepare the system to be restarted. The system will halt all I/O requests and save the data to the disks before restarting. The restart process may take several minutes to complete.

3.9.2 System Shutdown

NOTE : Use **SHUTDOWN** whenever you need to power down a controller for maintenance. **SHUTDOWN** flushes any data left in the cache and prepares the SFA10000 for an orderly shutdown.

The **SHUTDOWN CONTROLLER [LOCAL|REMOTE|0|1]** command shuts down the specified controller.

The **SHUTDOWN SUBSYSTEM** command shuts down both controllers.

If you need to power down the SFA10000, use **SHUTDOWN** prior to shutting off the power. This will cause the SFA10000 to immediately flush its cache, abort all initialization and rebuild operations, and proceed with an orderly shutdown.

All hosts actively using the SFA10000 should be safely shutdown and all users logged out before using this command. The SFA10000 will halt all I/O requests and save the data to the disks. The system can be safely turned off after using this command. Use the Power button on the UPS unit to turn off power to the controller.

Once shut down is complete, all power supplies must be switched off or unplugged. Use the Power button on the UPS unit to remove power to the controller. Power must be removed from the system for at least 10 seconds before it will start up again.

For SFA10000E, the Virtual Machines are automatically shut down whenever the **SHUTDOWN SUBSYSTEM** command is issued.

3.10 Performance Management

3.10.1 Cache Coherency

By default, the SFA10000 runs in couplet mode, where both controllers are running simultaneously, communicating through internal Ethernet and InfiniBand connections. This means that the SFA10000 is always running with cache coherency enabled.

The cache settings are configurable for each storage pool. Available keywords are:

FULL_STRIPE_WRITE_CACHING Specify full stripe write caching
MIRROREDSpecify cache data mirroring
READ_AHEAD_CACHING Specify read ahead caching
WRITE_BACK_CACHING Specify write back caching

To display the current settings, enter the command **SHOW POOL=<id> ALL_ATTRIBUTES**.

To change the settings, use these commands:

```
SET POOL=<id> FULL_STRIPE_WRITE_CACHING=[TRUE|FALSE]
SET POOL=<id> MIRRORED=[TRUE|FALSE]
SET POOL=<id> READ_AHEAD_CACHING=[TRUE|FALSE]
SET POOL=<id> WRITE_BACK_CACHING=[TRUE|FALSE]
```

NOTE : Turning off mirroring may result in data integrity issues.

3.10.1.1 Cache Protection

The SFA10000 UPS serves as a battery, which allows the controller to flush the contents of cache to non-volatile storage (internal disk) during a power event. When cache has been completely flushed, the controller will shut itself down. When power is restored, the flushed cache is replayed back from the non-volatile storage and committed to disk.

Mirrored Write Back Cache (MWBC) provides a mechanism of cache protection by copying, or mirroring, the contents of cache from one singlet to the peer singlet (and vice-versa). In the event of a singlet failure, the copy of cache sitting in the surviving peer can be committed to disk by the surviving peer. Although there is a latency involved in the cache mirroring process, it is recommended as additional protection of data.

3.10.2 Single Controller Write Back Cache Policy

The SFA10000 can be set to automatically switch to write-thru mode upon loss of one of the controllers in the couplet configuration. Loss of a controller is defined as a controller shutting down for any reason, such as crash, normal shutdown, and hardware failure.

To change the settings, enter the command:

```
SET SUBSYSTEM WRITE_BACK_DISABLE=[TRUE|FALSE]
```

The default setting is **FALSE** which allows automatic switching to write-thru mode upon loss of one controller.

3.10.3 Right Side I/O

Access to virtual disks uses the notion of a preferred home or "homed" path. If a path to the VD utilizes the peer controller, additional latency is incurred for every I/O as it must be processed first by the "non-homed" controller and forwarded over the same communication path used by Mirrored Write Back Cache (Figure 91).

Figure 91. Path to VD Utilizes the Peer Controller

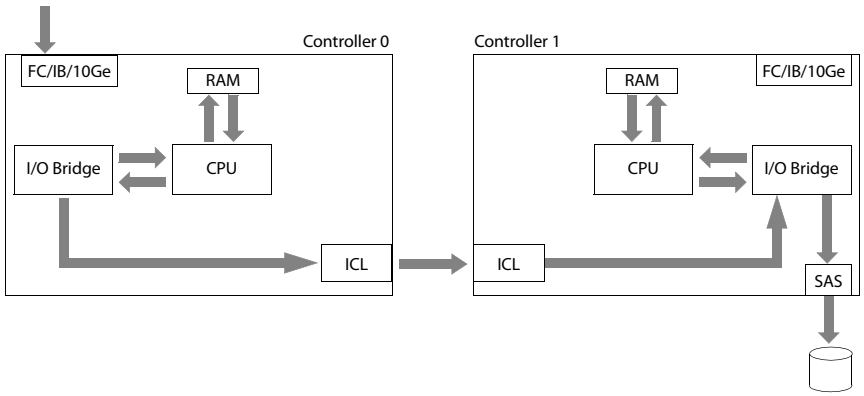

Properly configured and tested host side multi-path drivers will ensure right side I/O will occur. If poor performance is encountered, especially upon initial configuration, then verifying and correcting your primary path to disks is essential.

Verify by reviewing the virtual disk counters, enter the command (Figure 92):

```
SHOW VD * COUNTERS
```

NOTE : You need to enter this command three times—first time to initiate the counters, second time to display the results, third time to reset the counters.

Controller 0's results are displayed on the left and controller 1's results are displayed on the right. In Figure 92, if VD 102 is mastered by controller 0, all of the I/O is passing through controller 1, thus doing wrong side I/O.

Figure 92. Show Virtual Disk Counters Example Screen

```
RAID[0]$ show vd * counters

Virtual disk Counters: Elapsed time = 12.181 seconds

  Idx  IOs/sec   KiB/sec   KiB/IO Fwd IO/s Fwd KiB/s| IOs/sec   KiB/sec   KiB/IO Fwd IO/s Fwd KiB/s|
-------------------------------------------------------------------------------------------------------
  102        0         0        0       0         0 |    6580    109243       68       0         0 |
  103        0         0        0       0         0 |    7005    116304       68       0         0 |
  104        0         0        0       0         0 |    5158    186392      148       0         0 |
  105        0         0        0       0         0 |    2825    102083      148       0         0 |
  106        0         0        0       0         0 |    3481    210768      248       0         0 |
  107        0         0        0       0         0 |    3641    220500      248       0         0 |
  108        0         0        0       0         0 |       0         0        0       0         0 |
  109        0         0        0       0         0 |       0         0        0       0         0 |
  110        0         0        0       0         0 |       0         0        0       0         0 |
  111        0         0        0       0         0 |       0         0        0       0         0 |
  112        0         0        0       0         0 |       0         0        0       0         0 |
  113        0         0        0       0         0 |       0         0        0       0         0 |
  114        0         0        0       0         0 |       0         0        0       0         0 |
  115        0         0        0       0         0 |       0         0        0       0         0 |
  116        0         0        0       0         0 |       0         0        0       0         0 |
  117        0         0        0       0         0 |       0         0        0       0         0 |
  118        0         0        0       0         0 |       0         0        0       0         0 |
  119        0         0        0       0         0 |       0         0        0       0         0 |
```

3.10.4 Background Job Priority

Initialization and rebuild operations are background processes and their rates can be adjusted to minimize their impact on system performance.

To display the list of current jobs, enter command:

SHOW JOB *

The type of job and percentage of completion are displayed (Figure 93).

Figure 93. Show Job Example Screen (1)

```
**************************
*     Background Jobs    *
**************************
                                                    |Fraction|
 Idx|Type           |Target   (Sub)    |State      |Complete|Priority|Status   |Time
-------------------------------------------------------------------------------------
   0 FULL REBUILD    POOL:0   (PD:135 ) RUNNING        53%      80%              NA
   1 VERIFY ONCE     POOL:1   (NA)       RUNNING         0%      99%              NA
   2 VERIFY ONCE     POOL:2   (NA)       RUNNING         0%      99%              NA

Total Background Jobs: 3
```

To display the details of a job (Figure 94), enter command:

SHOW JOB <job-idx> ALL_ATTRIBUTES

Figure 94. Show Job Example Screen (2)

```
**************************
*     Background Jobs    *
**************************
  OID:                0x2b220000 (Index:0    )
  Target:             POOL:0x199b0000 (Index:0    )
  Sub-Target:         PD:0x21310087 (Index:135  )
  Type:               REBUILD
  State:              RUNNING
  Completion Status:  UNKNOWN
  Priority:           80%
  Fraction Complete:  52%
  Time:               NA
```

You may specify the amount of system resources that should be devoted to a background job. The higher its priority value, the faster the background job will run and the more the background job will impact client I/O performance.

To set the job priority, enter the command:

SET JOB=<id> PRIORITY=<priority>

where <priority> is a number between 1 and 99. Note that PRIORITY is not a percentage or a mathematical fraction of the available resources. For example, two background jobs with priority values of 50 will not use 100% of the system resources. The SFA10000 may or may not limit the number of background jobs to keep the total of their priorities below 100.

3.10.4.1 Pause/Resume a Job

You may pause a job at any time using the **PAUSE JOB=<idx>** command (Figure 95).

Figure 95. Pause a Job

```
RAID[0]$ pause job 0
JOB 0 OID=0x2b010000 paused with STATUS='Janus Success' (0x0)

RAID[0]$ show job 0 all
  OID:               0x2b010000 (Index:0)
  Target:            POOL:0x1a520000 (Index:0)
  Sub-Target:        NA
  Type:              REBUILD
  Status:            PAUSED
  Completion Status: UNKNOWN
  Priority:          50
  Fraction Complete: 11%
  Time:              NA
```

To resume the job, enter the command: **RESUME JOB=<id>** (Figure 96)

Figure 96. Resume a Job

```
RAID[0]$ resume job 0
JOB 0 OID=0x2b010000 resumed with STATUS='Janus Success' (0x0)

RAID[0]$ show job 0 all
  OID:               0x2b010000 (Index:0)
  Target:            POOL:0x1a520000 (Index:0)
  Sub-Target:        NA
  Type:              REBUILD
  Status:            RUNNING
  Completion Status: UNKNOWN
  Priority:          50
  Fraction Complete: 11%
  Time:              NA
```

3.10.5 Rebuild Policy Priority

Each storage pool has its own rebuild policy settings. You may specify the priority values for full rebuild and partial (fractional) rebuild policies. To display the current settings, enter the command **SHOW POOL=<id> ALL_ATTRIBUTES** (Figure 97).

Figure 97. Show Pool Information

```
***************
*   Pool(s)   *
***************
  Index:                  3
  OID:                    0x19b40003
  Type:                   STORAGE
  Name:                   raid 5 set
  Chunk Size:             64KB (0x80 blocks)
  Block Size:             512
  RAID Type:              RAID5
  Free Raid5 Capacity:    14576 GB
  Max VD Size:            14576 GB
  Total Capacity:         18240 GB
  UUID:                   0x00
  Global Spare Pool:      pool-6
  DiskTimeout(FRT):       10 minutes
  Init Policy:            ALLOW_IO
  Init Priority:          50%
  Full Rebuild Priority:        80%
  Fractional Rebuild Priority:  90%
  Sparing Policy:         AUTOMATIC
    .
    .
```

Full Rebuild Priority

To change the full rebuild priority value (Figure 98), enter the command:

SET POOL=<id> REBUILD_FULL_POLICY=<priority>

where <priority> is an integer in the range of 1 to 99. The default value specifies a priority of 80%.

Figure 98. Set Full Rebuild Priority

```
RAID[0]$ set pool 3 rebuild_full_policy=70

POOL 3 OID=0x19b40003 attributes set with STATUS='Success' (0x0)
```

Partial Rebuild Priority

To change the partial rebuild priority value (Figure 99), enter the command:

SET POOL=<id> REBUILD_PARTIAL_POLICY=<priority>

where <priority> is an integer in the range of 1 to 99. The default value specifies a priority of 90%.

Figure 99. Set Partial Rebuild Priority

```
RAID[0]$ set pool 3 rebuild_partial_policy=80

POOL 3 OID=0x19b40003 attributes set with STATUS='Success' (0x0)
```

3.11 Firmware Update Management

DataDirect Networks periodically releases firmware updates to enhance features of the products. Please contact our Technical Support Department to obtain the latest firmware files.

Before performing any updates, please note the following:

- The **UPDATE ENCLOSURE** command starts a background operation so that the CLUI is free to do other operations during the **UPDATE**. The status that comes back from the **UPDATE ENCLOSURE** command only indicates whether the background operation was successfully started. **SHOW ENCLOSURE <id> UPDATE_FIRMWARE_PROGRESS** should be used immediately after the command and then subsequently periodically to check the status of the **UPDATE**.
- When you SCP the Consolidated Download File (.DDN) to the firmware directory, it must be copied to the PRIMARY controller. Since the PRIMARY may change, it is probably best to copy it to BOTH controllers. If **SHOW ENCLOSURE <id> UPDATE_FIRMWARE_PROGRESS** reports status JS_ES_DL_UCODE_OPEN_FAILED, the most likely cause is that the file was not copied to the PRIMARY controller.
- When the SFA10000 is busy rebuilding, initializing, or doing host I/O, **UPDATE ENCLOSURE** operations may fail because the enclosure is too busy to service the UPDATE IOs in a timely fashion. Therefore, make sure that there are no rebuilds, initializations, or host I/O operations being done during the **UPDATE ENCLOSURE** operation.
- In order to cause the new enclosure firmware to start being used you must power cycle the enclosure. In the SFA10000, it is best to do this by doing a **SHUTDOWN SUBSYSTEM** followed by power-cycling the enclosure. This is also required to recover from an **UPDATE ENCLOSURE** failure. These power cycles are required.
- The **UPDATE PD** command starts a background operation so that the CLUI is free to do other operations during the **UPDATE**. The status that comes back from the **UPDATE PD** command only indicates whether the background operation was successfully started. **SHOW PD * UPDATE_FIRMWARE_PROGRESS** should be used immediately after the command and then subsequently periodically to check the status of the **UPDATE**.
- When you SCP the Consolidated Download File (.DDN) to the firmware directory, it must be copied to the PRIMARY controller. Since the PRIMARY may change, it is probably best to copy it to BOTH controllers. If **SHOW PD * UPDATE_FIRMWARE_PROGRESS** reports status JS_ES_PD_DL_UCODE_OPEN_FAILED for a physical disk, the most likely cause is that the file was not copied to the PRIMARY controller.
- Online physical disk firmware update is not currently supported. Therefore, make sure that there are no rebuilds, initializations, or host IO operations being done during the **UPDATE PD** operation. This is enforced by making sure the user has issued the **SET SUBSYSTEM OFFLINE** command before allowing an **UPDATE PD** command.
- In order to cause the new physical disk firmware to start being used you must power cycle the enclosure in which the physical disk is installed. In the SFA10000, it is best to do this by doing a **SHUTDOWN SUBSYSTEM** followed by power-cycling the enclosure.

3.11.1 Displaying Current Firmware Version

The **SHOW CONTROLLER <id> ALL_ATTRIBUTES** and **SHOW ENCLOSURE <id> ALL_ATTRIBUTES** commands display version information of the SFA10000's hardware and firmware (Figure 100).

Figure 100. show Controller Information Screen

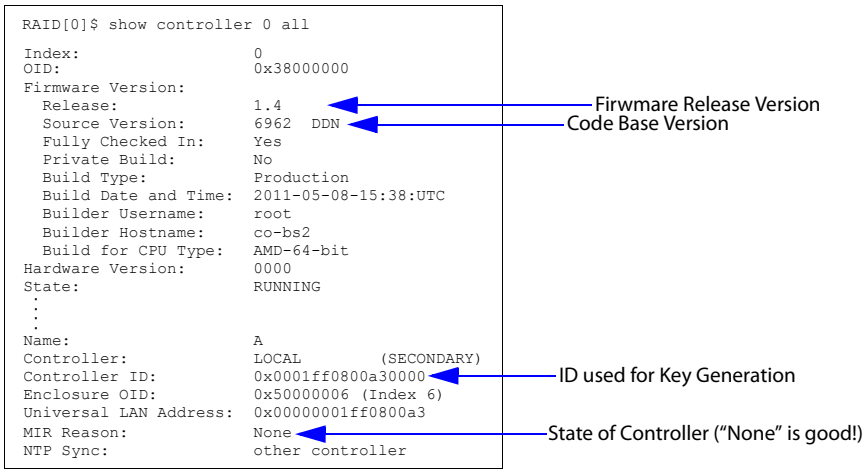

3.11.2 Controller Firmware Update Procedure

The firmware update procedure described here only applies to upgrades from firmware version 1.0.3.3.xxxx or later to a new version. Upgrades from an older version of firmware must be done by a trained technician.

> **NOTE :** The SFA10000 now supports an upgrade of the firmware while the system is still online. However, you must correctly follow the upgrade instructions to perform a successful online upgrade.

If you are upgrading from v1.0.3.3.xxxx or later, there are two methods that can be utilized:

- Copy the new firmware image over to the controller via the network instead of using a USB flash disk for Linux.
- Copy the new firmware image over to the controller via the network for Windows users.

> **NOTE :** In the examples, the filename of the firmware image is **ddn-flash-2812-opt.tgz** and the IP address of the controller is **10.32.31.240**. Replace these parameters with the filename and IP address appropriate for your installation.

3.11.2.1 Linux Environment-Firmware Update from the Network

Follow these steps to update the firmware:

1. Copy the new firmware to both controllers using an scp or sftp client. The user name is **`firmware`** and the password is **`Firmware`** (Note that entries are case-sensitive).
   For example:
   `scp ddn-flash-2812-opt.tgz firmware@10.32.31.240:`
2. At the CLUI prompt, enter command:
   `UPDATE_FIRMWARE CONTROLLER=LOCAL FILE="<file-specification>"`

   For example:
   `UPDATE_FIRMWARE CONTROLLER=LOCAL FILE="ddn-flash-2812-opt.tgz"`
   (Note that the file name must be enclosed with double quotation marks.)
3. Upon completion of reboot, login and enter the command:
   `SHOW CONTROLLER LOCAL ALL`
   Verify that the firmware version is correct.

   NOTE : The pools may indicate there is a fault. Issue the command **`SHOW POOLS *`** to check the details of the pool. The pools will go into WR thru state on initial startup to allow the batteries the opportunity to charge. The WR Thru mode is necessary to allow for data integrity until battery redundancy on each controller is completed. If there is a battery issue, the system will indicate replacement is required and all VDs will remain in WR Thru mode until the issue is resolved.

   For example:
   ```
   CLUI unable to get RAID SUBSYSTEM NAME. STATUS='MIR:Firmware Version Mismatch'
   (0x30003f8)
   ```

   NOTE : The firmware will be in MIR, firmware version mismatch state for the first controller that is upgraded. You must ensure that you have a CLUI prompt from the newly upgraded controller before moving to Step 4.
4. Upgrade the second controller by repeating Steps 2-3 above on the other controller.

3.11.2.2 Windows Environment-Firmware Update from the Network

This method requires a SFTP or SCP client, and involves two steps – copying the firmware image to both controllers and launching the update.

1. Copy the new firmware to the controller using an scp or sftp client. You must use an application that supports SFTP or SCP. WinSCP, a free SFTP and SCP client, is used in this example.
2. Launch WinSCP or other SFTP/SCP client. For hostname, use the IP address of your controller. user name is **`firmware`** and password is **`Firmware`**. Both user name and password are case-sensitive.
3. Select Login .

   At the login screen (Figure 101), use the user name **`admin`** with the password **`password`** to log into the system. User name and passwords are care sensitive.

Figure 101. Login Screen

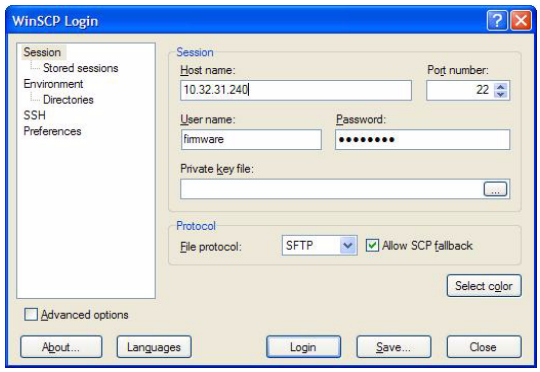

4. Find the image file and drag this to the destination window and select copy. Copy operation commences.

To launch the update:

1. Login to the controller via an SSH session with the user name **user** and password **user**.
2. At the CLI prompt, enter:
   update_firmware controller=<controller-id> FILE="<file-specification>"

   For example:
   update_firmware controller=LOCAL FILE= "ddn-flash-2812-opt.tgz"

   NOTE : The quotes around the filename of the firmware image are required.

3. Upon completion of reboot, login and enter the command:
   SHOW CONTROLLER LOCAL ALL
   Verify that the firmware version is correct.

   NOTE : The pools may indicate there is a fault. Issue the command **SHOW POOLS *** to check the details of the pool. The pools will go into WR thru state on initial startup to allow the batteries the opportunity to charge. The WR Thru mode is necessary to allow for data integrity until battery redundancy on each controller is completed. If there is a battery issue, the system will indicate replacement is required and all VDs will remain in WR Thru mode until the issue is resolved.

   For example:
   ```
   CLUI unable to get RAID SUBSYSTEM NAME. STATUS='MIR:Firmware Version Mismatch'
   (0x30003f8)
   ```

   NOTE : The firmware will be in MIR, firmware version mismatch state for the first controller that is upgraded. You must ensure that you have a CLUI prompt from the newly upgraded controller before moving to Step 4.

4. Upgrade the second controller by repeating Steps 2-3 above on the other controller.

3.11.3 Disk Enclosure Firmware Update Procedure

NOTE : This is an offline upgrade. Any application/host should be shutdown or quiesced prior to performing this operation.

Known Issues

When upgrading the enclosure firmware, any errors experienced while downloading the firmware will not be realized until the operation is complete and the shelf is power cycled. At that time, any errors are displayed and it may be necessary to attempt the upgrade a second time.

It is required that you shutdown one controller prior to upgrading the enclosure firmware. Use the CLUI command: **`SHUTDOWN CONTROLLER REMOTE`**.

3.11.3.1 Enclosure File Definitions

To determine which file(s) you need to upload you will need to determine the enclosure types you have. To do this, issue the CLUI command: **`SHOW ENCLOSURE *`**

This will display the enclosure information in a table format.

```
************************
*     Enclosure(s)     *
************************
Idx|Type      |Logical ID           |Vendor ID|Product ID |Revision|FW Version
-------------------------------------------------------------------------------
  0 CONTROLLER 0x0001ff0801c90000     DDN       SFA10000     0000     1.4
  1 DISK       0x50001ff210030000     DDN       SS7000       0402     04.02.01
  2 DISK       0x50001ff210006000     DDN       SS7000       0402     04.02.01
  3 DISK       0x50001ff21000a000     DDN       SS7000       0402     04.02.01
  4 DISK       0x50001ff210008000     DDN       SS7000       0402     04.02.01
  5 DISK       0x50001ff21002a000     DDN       SS7000       0402     04.02.01
  6 CONTROLLER 0x0001ff0802810000     DDN       SFA10000     0000     1.4

Total Enclosures: 7
```

The naming convention of the files match up accordingly:

- SFA10000 controller - SFA10000_<FW Version>.DDN
- SS6000 - SS6000_<FW Version>.DDN
- SS7000 - SS7000_<FW Version>.DDN
- SS2460 - SS2460_<FW Version>.DDN

Before starting this procedure, verify that a power cycle or any perturbations to the system is NOT required for the next 2 hours.

There are two (2) parts to update the enclosure firmware:

1. Upload the enclosure firmware to the controller using either Linux Environment or Windows.
2. Update Disk Shelf using CLUI commands.

3.11.3.2 Linux Environment-Firmware Upload

Obtain the firmware from DDN. Firmware must be running on the controller and must be connected to the enclosure to be upgraded.

1. Use the secured copy program (scp) to transfer the file to the expander with the command: **`scp <firmware file name> firmware@<ip_address_of_SFA10000>:`**
   To determine the file to upload, refer to Section 3.11.3.1, "Enclosure File Definitions".

   NOTE : Ensure the colon (:) is at the end of the above command.

2. Enter the password **`Firmware`**. (Note that entries are case-sensitive).

The file will be copied to the controller.

3.11.3.3 Windows Environment-Firmware Upload

Obtain the firmware from DDN. Firmware must be running on the controller and must be connected to the enclosure to be upgraded.

1. Use the putty secured copy program (pscp) to transfer the file to the expander with the command: **`pscp <firmware file name> firmware@<ip_address_of_SFA10000>:`**

   NOTE : Ensure the colon (:) is at the end of the above command.

2. Enter the password **`Firmware`**. (Note that entries are case-sensitive).

The file will be copied to the controller.

3.11.3.4 Disk Shelf Upgrade

Once the file is copied from either procedure above (Linux or Windows), follow the steps below to complete the upgrade:

1. Log into the controller.
   User name is **`user`**. Password is **`user`**.

2. One of the controllers needs to be shut down after the subsystem is offline. At the CLUI prompt, enter command: **`SHUTDOWN CONTROLLER LOCAL`**

   This will shut down the local controller. You will then need to log into the other controller and issue the upgrade command as shown in the next step. This will need to be done for each of the disk enclosures.

3. At the CLUI prompt, enter the command:
   `UPDATE ENCLOSURE <enclosure num> FILE=<file name uploaded>`

   To determine the enclosure number in a system issue the command: **`SHOW ENCLOSURE *`**. This will display the list of enclosures connected to the system as shown below.

```
************************
*     Enclosure(s)     *
************************
Idx|Type       |Logical ID           |Vendor ID|Product ID |Revision|FW Version
-------------------------------------------------------------------------------
  0 CONTROLLER 0x0001ff0801c90000     DDN       SFA10000     0000     1.4.0
  1 DISK       0x50001ff210030000     DDN       SS7000       0402     04.02.01
  2 DISK       0x50001ff210006000     DDN       SS7000       0402     04.02.01
  3 DISK       0x50001ff21000a000     DDN       SS7000       0402     04.02.01
  4 DISK       0x50001ff210008000     DDN       SS7000       0402     04.02.01
  5 DISK       0x50001ff21002a000     DDN       SS7000       0402     04.02.01
  6 CONTROLLER 0x0001ff0802810000     DDN       SFA10000     0000     1.4.0

Total Enclosures: 7
```

The upgrade will take approximately 1 hour for each enclosure.

4. To monitor the upgrade progress, enter the command:
SHOW ENCLOSURE 1 UPDATE_FIRMWARE_PROGRESS

   For example:

```
RAID[0]$ show enclosure 1 update_firmware_progress
ES download progress for enclosure 1.
  CDF file name           /tmp/janus_update//SFA_D02_011.DDN
  CDF is for Vendor ID    DDN
  CDF is for Product ID   SFA OS
  CDF package version     D02.011
  CDF release date        01-Mar-2011
  Download is in progress and 15 percent complete.
  Download consists of 2 images.
  Image 2 is in progress and 13 percent complete.
```

> **NOTE :** An alternate command to monitor the upgrade progress, **SHOW ENCLOSURE 1 UPDATE_FIRMWARE_PROGRESS+**. The "+" at the end of the command will automatically reissue the command until a key is depressed. You can watch the upgrade status without having to re-enter the command.

5. Verify if the update is completed, enter the command:
SHOW ENCLOSURE 1 UPDATE_FIRMWARE_PROGRESS

   For example:

```
RAID[0]$ show enclosure 1 update_firmware_progress
ES download progress for enclosure 1.
  Download not in progress -- last download completed successfully.
```

6. When the download is completed, you must power-cycle the enclosure. **Prior to powering off the enclosure, issue the SHUTDOWN SUBSYSTEM command**.

> **NOTE :** Make sure the power is left off for at least 30 seconds after the power down or the controller may not boot up correctly the first time.

7. Verify that the upgrade was successful by using the command: **SHOW EXPANDER**

> **NOTE :** The expanders are Sub Index 3 thru 10. Ensure that the Firmware version and the Init String version are identical. If the **Firmware version** field indicates "VARIOUS", this indicates that the expanders, i.e. DEMs, have not accepted the update. You will need to power cycle the enclosure again to try and clear this condition. If the "VARIOUS" state still exists after the second power cycle, you will need to re-run the upgrade procedure.

For example:

```
RAID[0]$ show expander 1 * all
Sub Index:              1
Sub OID:                0x78000001
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               1
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Location:               IOM1
Firmware version:       0078

Sub Index:              2
Sub OID:                0x78000002
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               2
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Location:               IOM1

Sub Index:              3
Sub OID:                0x78000003
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               3
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Location:               DEM1
Part number:            TCA-00300-01-A
Serial number:          MXSCI00089HVD12B
Firmware version:       D02.011
Init string version:    D02.011
FPGA version:           7

Sub Index:              4
Sub OID:                0x78000004
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               4
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Location:               DEM2
Part number:            TCA-00300-01-A
Serial number:          MXSCI00089HVD110
Firmware version:       D02.011
Init string version:    D02.011
FPGA version:           7

Sub Index:              5
Sub OID:                0x78000005
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               5
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Location:               DEM3
Part number:            TCA-00300-01-A
Serial number:          MXSCI00085QVD2CB
Firmware version:       D02.011
Init string version:    D02.011
FPGA version:           7

Sub Index:              6
Sub OID:                0x78000006
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               6
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Location:               DEM4
Part number:            TCA-00300-01-A
Serial number:          MXSCI00086QVD109
Firmware version:       D02.011
Init string version:    D02.011
FPGA version:           7

Sub Index:              7
Sub OID:                0x78000007
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               7
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Location:               DEM5
Part number:            TCA-00300-01-A
Serial number:          MXSCI000893VD19D
Firmware version:       D02.011
Init string version:    D02.011
FPGA version:           7

Sub Index:              8
Sub OID:                0x78000008
Enclosure Index:        1
```

```
Enclosure OID:          0x50000001
Position:               8
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Location:               DEM6
Part number:            TCA-00300-01-A
Serial number:          MXSCI00085LVD0A5
Firmware version:       D02.011
Init string version:    D02.011
FPGA version:           7

Sub Index:              9
Sub OID:                0x78000009
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               9
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Location:               DEM7
Part number:            TCA-00300-01-A
Serial number:          MXSCI00089HVD11A
Firmware version:       D02.011
Init string version:    D02.011
FPGA version:           7

Sub Index:              10
Sub OID:                0x7800000a
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               10
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Location:               DEM8
Part number:            TCA-00300-01-A
Serial number:          MXSCI00089HVD12C
Firmware version:       D02.011
Init string version:    D02.011
FPGA version:           7


Total Expanders: 10
```

3.11.4 Physical Disk Firmware Update Procedure

There are two (2) parts to update the physical disk firmware:

1. Upload the physical disk firmware to the SFA10000 using either Linux Environment or Windows.
2. Update the Physical Disk using CLUI commands.

3.11.4.1 Linux Environment-Firmware Upload

Obtain the firmware from DDN.

1. Use the secured copy program (scp) to transfer the file to the controller with the command: **`scp <firmware file name> firmware@<ip_address_of_SFA10000>:`**

   NOTE : Ensure the colon (:) is at the end of the above command.

2. Enter the password (case-sensitive) **`Firmware`**.

The file is copied to the controller.

3.11.4.2 Windows Environment-Firmware Upload

Obtain the firmware from DDN.

1. Use the putty secured copy program (pscp) to transfer the file to the controller with the command:
   pscp <firmware file name> firmware@<ip_address_of_SFA10000>:

   NOTE : Ensure the colon (:) is at the end of the above command.

2. Enter the password (case-sensitive) **Firmware**.

The file is copied to the controller.

3.11.4.3 Physical Disk Upgrade

Once the file is copied from either procedure above (Linux or Windows), follow the steps below to complete the upgrade:

1. Log onto the controller.
   User name is **user**. Password is **user**.

2. At the CLUI prompt, enter command:
   SET SUBSYSTEM OFFLINE

   NOTE : The subsystem must be offline before proceeding with the remaining steps.

3. At the CLUI prompt, enter command:
   UPDATE PD * FILE=<file name uploaded>

   The firmware will be updated on any physical disks that match the Vendor Id, Product Id, and Drive Type (SAS/SATA) in the file.

   The upgrade time varies, but typically takes less than a minute.

   For example:

```
RAID[0]$ update pd * file="WD2002FYPS_05D07.DDN"
A request has been made to update physical disk firmware. No attempt will be made to preserve
redundancy if the physical disk resides in a pool. If the update to the physical disk
encounters catastrophic failure, then this may result in lost data if the pool loses all
of its redundancy.
It is recommended that a backup is done prior to this operation.
Are you sure you want to update physical disk firmware [Yes]?
PHYSICAL DISK 65535 OID=0xffffffff firmware update started STATUS='An asynchronous command
has been started' (0x3000068)
The controller has initiated the download operation which may take up to several minutes.
Check the progress and status of the download operation with SHOW PD n
UPDATE_FIRMWARE_PROGRESS. Additional information may be found in the event log.
```

4. To monitor the upgrade progress, enter command:
   SHOW PD * UPDATE_FIRMWARE_PROGRESS

   For example (non-matching drives omitted in output for clarity):

```
RAID[0]$ show pd * update_firmware_progress
Encl|Slot| Vendor  |     Product ID    |Type| Revision| Serial Number  | Pool| Idx|      WWN        |% Complete|   Status
-------------------------------------------------------------------------------------------------------------------------------
1      5   HITACHI    HUS153014VLS300    SAS     A410    JFVWMZNC5301      UNAS  271  50014ee2032d2555   56%         In Progress
1      6   HITACHI    HUS153014VLS300    SAS     A410    J4V1YYAA5301      UNAS  272  50014ee2032d196c   56%         In Progress
1     17   HITACHI    HUS153014VLS300    SAS     A410    JFVWL5YC5301      UNAS  274  50014ee2032ca9a9   56%         In Progress
1     18   HITACHI    HUS153014VLS300    SAS     A410    JFVSU5PC5301      UNAS  269  50014ee25882b499   56%         In Progress
```

5. When the upgrade has completed, to verify that the download succeeded, enter command:
   SHOW PD * UPDATE_FIRMWARE_PROGRESS

   For example (non-matching drives omitted in output for clarity):

```
RAID[0]$ show pd * update_firmware_progress
Encl|Slot| Vendor |     Product ID  |Type| Revision|   Serial Number| Pool|Idx|    WWN          |% Complete|Status
-------------------------------------------------------------------------------------------------------------------------
1     5   HITACHI   HUS153014VLS300  SAS     A410    JFVWMZNC5301     UNAS  271  50014ee2032d2555    Done    JS_GBL_SUCCESS (Success)
1     6   HITACHI   HUS153014VLS300  SAS     A410    J4V1YYAA5301     UNAS  272  50014ee2032d196c    Done    JS_GBL_SUCCESS (Success)
1    17   HITACHI   HUS153014VLS300  SAS     A410    JFVWL5YC5301     UNAS  274  50014ee2032ca9a9    Done    JS_GBL_SUCCESS (Success)
1    18   HITACHI   HUS153014VLS300  SAS     A410    JFVSU5PC5301     UNAS  269  50014ee25882b499    Done    JS_GBL_SUCCESS (Success)
```

6. Shutdown the subsystem, enter command:
   SHUTDOWN SUBSYSTEM

7. Power-cycle the enclosures containing the updated physical disks.

8. After powering up the disk enclosures and controller, verify the upgrade was successful, enter command:
   SHOW PD *

   The firmware version of the updated physical disks should indicate the new version.

   For example (non-matching drives omitted in output for clarity):

```
RAID[0]$ show pd *
                                                                                                |Health|
Encl|Slot| Vendor |    Product ID   |Type| Cap GB | RPM|Revision| Serial Number  |Pool| State|Idx | Stat
-------------------------------------------------------------------------------------------------------
1    5     HITACHI HUS153014VLS300   SAS     1863  5.4K A410      JFVWMZNC5301    UNAS  GOOD  279   READ
1    6     HITACHI HUS153014VLS300   SAS     1863  5.4K A410      J4V1YYAA5301    UNAS  GOOD  280   READ
1   17     HITACHI HUS153014VLS300   SAS     1863  5.4K A410      JFVWL5YC5301    UNAS  GOOD  282   READ
1   18     HITACHI HUS153014VLS300   SAS     1863  5.4K A410      JFVSU5PC5301    UNAS  GOOD  277   READ
```

3.12 The System Logs

3.12.1 Displaying Event Logs

To display the event log starting at the start sequence number in ascending sequence number order (Figure 102), enter command:

SHOW CONTROLLER <id> LOG ASCEND

Figure 102. Show Controller Log Example Screen (1)

```
RAID[0]$ show controller 0 log ascend

000001 2010-06-04 01:03:26:6418393 G=0 S=0 T=1 RP=0 VP=1
LOG_ES_UPS_CLIENT_ATTRIBUTE_NAME  ES Controller enclosure 0x1ff0800a30000 UPS attribute 0.....

000002 2010-06-04 01:03:26:6418400 G=0 S=0 T=1 RP=0 VP=1
LOG_ES_UPS_CLIENT_ATTRIBUTE_CHANGED  ES Controller UPS attribute 0 changed to 37......

000003 2010-06-04 01:04:52:8320854 G=3 S=1 T=1 RP=0 VP=1
LOG_LOGDISK_ENABLE_RECEIVED_FROM_STATE LOG RECEIVED FROM STATE
```

To display the event log starting at the start sequence number in descending sequence number order (Figure 103), enter command:

SHOW CONTROLLER <id> LOG DESCEND

Figure 103. Show Controller Log Example Screen (2)

```
RAID[0]$ SHOW CONTROLLER 1 LOG DESCEND

000024 2010-02-11 05:08:48:7027390 G=0 S=0 T=1 RP=0 VP=1 LOG_JOI_TIME_SET JOI TIME
WAS SET BY AN ADMINISTRATOR AT 2010-2-11 5:8:48; NEARBY LOG ENTRIES MAY APPEAR OUT
OF TIME ORDER. OFFSET= 0X1C987C765CD3B2B.

000023 2010-02-06 04:49:20:5069068 G=3 S=1 T=1 RP=0 VP=1 LOG_LOGDISK_ENABLE_RECEIVED
_FROM_STATE LOGRECEIVED FROM STATE

000022 2010-02-06 04:49:20:4952631 G=0 S=0 T=1 RP=0 VP=1 LOG_JOI_TIME_SET JOI TIME
WAS SET BY AN ADMINISTRATOR AT 2010-2-6 4:49:20; NEARBY LOG ENTRIES MAY APPEAR OUT
OF TIME ORDER. OFFSET= 0X1C987C7A302483D.
```

To display the event log starting at the specified start sequence number in ascending sequence number order (Figure 104), enter command:

SHOW CONTROLLER <id> LOG ASCEND START_SEQUENCE <start-sequence-number>

Figure 104. Show Controller Log Example Screen (3)

```
RAID[0]$ SHOW CONTROLLER 1 LOG ASCEND START_SEQUENCE 20

000021 2010-02-06 04:49:20:4731906 G=4 S=2 T=1 RP=0 VP=1 LOG_ST_MIR_STATE STATE
 MIR STATE STATE:000A

000022 2010-02-06 04:49:20:4952631 G=0 S=0 T=1 RP=0 VP=1 LOG_JOI_TIME_SET JOI TIME
WAS SET BY AN ADMINISTRATOR AT 2010-2-6 4:49:20; NEARBY LOG ENTRIES MAY APPEAR
OUT OF TIME ORDER.OFFSET = 0X1C987C7A302483D.

000023 2010-02-06 04:49:20:5069068 G=3 S=1 T=1 RP=0 VP=1 LOG_LOGDISK_ENABLE_
RECEIVED_FROM_STATE LOG RECEIVED FROM STATE

000024 2010-02-11 05:08:48:7027390 G=0 S=0 T=1 RP=0 VP=1 LOG_JOI_TIME_SET JOI TIME
WAS SET BY AN ADMINISTRATOR AT 2010-2-11 5:8:48; NEARBY LOG ENTRIES MAY APPEAR
OUT OF TIME ORDER. OFFSET  = 0X1C987C765CD3B2B.
```

3.12.2 Event Log Structure

Figure 105 below illustrates the structure of an event log.

The severity levels include:

0: Informational
1: Warning
2: Error
3: Fatal

Some additional parameters are:

ES: Enclosure WWN:Slot number (for example, 50001ff101ed0000:60)
IOC: SAS/SATA controller number IOC0 (for example, WWN:5000cca216ed8430)
RC: Reason Code (04 - Device Not Responding)
PHY: Expander PHY device is attached too (for example, PHY:08)
ASCQ: ASC/ASCQ SCSI sense data

Figure 105. Event Log Structure

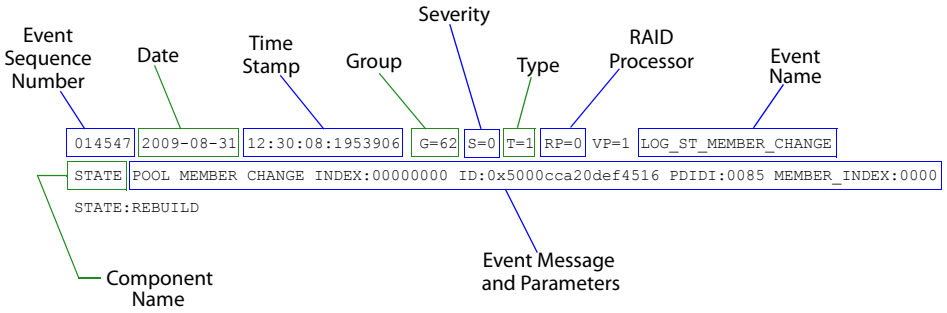

3.12.3 SFA OS Terms in the Log

- **STATE** – Implements all policy, implements the metadata store, handles all dual-controller issues, and controls all other modules
- **ES** (Enclosure Services) – Monitors all disk enclosures and its local controller enclosure
- **AMPD** and **MPI** – Implements the back-end (SAS/SATA) I/O controller drivers for access and discovery of Physical Disks and Enclosures
- **CM** (Cache Manager) – Implements read and write-back cache
- **RAID** – Implements RAID 5, RAID 6, RAID 1, and rebuilds
- **MAD** – Data transfer and parity calculations
- **RT** (Routing) – Routes VD request to the right CM
- **IOF** (I/O Forwarding) – Routes PD requests to the right AMPD
- **DUCK** – Dual-controller communication
- **JOI**, **JEX**, **JIPC**, **MIS**, **JTS** – Infrastructure, communication, and testing

3.13 Remote Management of SFA10000

The SFA10000 can be managed locally through the RS-232 interface, or remotely via SSH. The CLUI is the same regardless of the management interface (RS-232 or SSH).

3.13.1 Network Connection

Connect the Ethernet port on the back of the controllers to your Ethernet network (Figure 106). Then configure the network interface as described below.

Figure 106. Ethernet Connections to Your Network

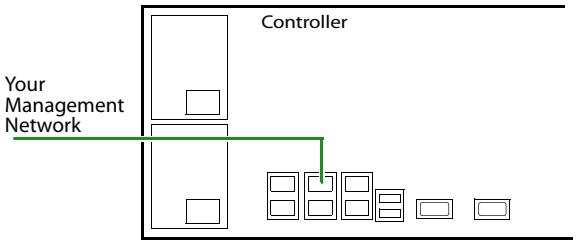

NOTE : Currently, the SFA10000 does not support network configuration protocols such as DHCP or BOOTP.

3.13.2 Display Network Interface Attributes

To display the current network interface settings, enter the command (Figure 107):
`UI SHOW NETWORK_INTERFACE [LOCAL|REMOTE|0|1] *`
where **`LOCAL`** gives you information on the controller that you are currently logged into.

Figure 107. Current Network Interface Settings Screen

```
RAID[0]$ ui show network_interface=local 0 *
Network device id 0
    address 10.32.31.31
    netmask 255.255.240.0
    gateway 10.32.16.2
```

3.13.3 Change Network Interface Settings

NOTE : Initial network interface settings must be configured using the serial interface. Refer to Section 2.7.7, "Configure Network Interface Settings" for instructions.

To change the network interface settings on the controller you are connected to, enter the command (Figure 108):
`UI SET NETWORK_INTERFACE=LOCAL 0 IP_ADDRESS=<ip_address> IP_MASK=<netmask> IP_GATEWAY=<gateway>`

Figure 108. Set Network Interface Example

```
RAID[0]$ ui set network_interface=local 0 ip_address=10.32.31.31 ip_mask=255.255.240.0 ip_gateway=10.32.16.2
NETWORK_INTERFACE 0 set with STATUS='Success' (0x0)
```

3.13.4 Logins

By default, the login name is **user** and its password is **user** (Figure 109). Both are case sensitive.

Only one SSH session is permitted at a time. Once a SSH session is initiated, the RS-232 console switches to a CLI sub-shell. The SSH client should be using port 22 with its local echoing function disabled.

Figure 109. SSH Login Screen

```
login as: user
user@10.23.23.16's password:

Linux (none) 2.6.25-ddn-016620-3 #2 SMP Wed Jan 14 10:38:28 MST 2009 x
DataDirect Networks Storage Scalar 6620(tm) J
```

Logout

To logout, enter command **QUIT**. For SSH connection, the current session will be disconnected.

3.13.5 Email and SNMP Notification Setup

3.13.5.1 Email Setup

There is only one logical EMAIL_AGENT for the subsystem. While each controller has an EMAIL_AGENT, the two are ideally managed as a single, logical EMAIL_AGENT in that all of their settable attribute values are shared.

Automatic emails will be sent as notification of a selected group of warning and error events that have occurred on the SFA10000. The recipient address is user-configurable and only one address is permitted per system.

To set the email address, enter the command (Figure 110):

UI SET EMAIL IP_ADDRESS=<mail server ip_address>
IP_PORT=<incoming mail server port> FROM= "<email address>"
TO="<email address>" SUBJECT="text"

To show all the email address attributes, use the **UI SHOW EMAIL ALL** command (Figure 110).

Figure 110. Email Setup Example Screen

```
RAID[0]$ ui set email ip_address=192.168.0.10
SNMP Agent set with STATUS='Success' (0x0)

RAID[0]$ ui set email ip_port=25
SNMP Agent set with STATUS='Success' (0x0)

RAID[0]$ ui set email to="jdoe@ddn.com"
SNMP Agent set with STATUS='Success' (0x0)

RAID[0]$ ui set email subject="DDN SFA10000 Event Notification"
SNMP Agent set with STATUS='Success' (0x0)


RAID[0]$ ui show email

  IP_ADDRESS=192.168.0.10
  IP_PORT=25
  FROM=jdoe@ddn.com
  TO=jdoe@ddn.com
  SUBJECT=DDN SFA10000 Event Notification
```

3.13.5.2 SNMP Setup

The Simple Network Management Protocol (SNMP) monitors network attached devices for conditions that warrant administrative attention. In Version 1.1, SNMP traps have been implemented to monitor critical and warning events. A management information base (MIB) has also been created to be used to provide inquiry objects and events to the user's monitoring application. The provided SNMP traps expose management data on the managed system in the areas of temperature senor, fans, power supplies, pools, and physical disks as well as a variety of real-time critical and error events.

Each controller has an SNMP_AGENT. When changing the SNMP settings, you must set the changes on each controller separately.

To set the SNMP trap agent on a controller, enter the command (Figure 111):
`UI SET SNMP IP_ADDRESS=<ip_address> COMMUNITY="<name>"`

To show all the SNMP trap agent attributes, use **`UI SHOW SNMP`** command (Figure 111).

Figure 111. SNMP Configuration Example Screen

```
RAID[0]$ ui set snmp ip_address=192.168.0.10
SNMP Agent set with STATUS='Success' (0x0)

RAID[0]$ ui set snmp community=public
SNMP Agent set with STATUS='Success' (0x0)

RAID[0]$ ui show snmp

SNMP Trap Agent Attributes

  IP_ADDRESS=192.168.0.10
  COMMUNITY=public
```

3.13.5.3 Inquiry Items and Events

Table 3, Table 4, and Table 5 below detail the additional inquiry items and events that will be trapped in the SNMP MIB and in the Email Agent.

Table 3. SNMP Inquiry Objects

SNMP Inquiry Objects	
Item	Returned Values
Temperature Sensor	Number of temperature sensors, list of temperature sensors. For each sensor : ID, Enclosure ID, Enclosure position, Status (normal, warning, critical).
Fans	Number of fans in the system, list of fans. For each fan: ID, Enclosure ID, Enclosure position, Status (healthy, failure).
Power Supplies	Number of power supplies in the system, list of power supplies. For each power supply: ID, Enclosure ID, Enclosure position, Status (healthy, failure).
Pools	Number of pools, list of pools. For each pool: OID, type (storage, spare, unassigned), access, number of members.
Physical Disks	List of disks, for each disk: WWN, enclosure number, slot, status (normal, failed, predicted failure, unknown)

Table 4. SNMP Only Events

SNMP ONLY Events	
Error	Description
LOG_AMPD_DEVICE_INIT_FAILED	A device failed to initialize.
LOG_AMPD_DSK_DEVICE_INIT_FAILED	DISK device initialization failure.
LOG_AMPD_ICL_DEVICE_INIT_FAILED	ICL device initialization failure.
LOG_AMPD_MPI_DSK_SCSI_SENSE_DATA	Disk data received.
LOG_AMPD_MPI_ICL_SCSI_SENSE_DATA	ICL SCSI sense data was received.
LOG_AMPD_MPI_SCSI_SENSE_DATA_EVT	SCSI sense data was received.
LOG_AMPD_MPI_SES_SCSI_SENSE_DATA	SES SCSI sense data was received.
LOG_AMPD_SES_DEVICE_INIT_FAILED	SES device initialization failure.
LOG_ES_BATTERY_FAILURE_PREDICTED	UPS for the controller is predicted to fail and should be replaced.
LOG_ES_COOLING_ELEMENT_INSERTED	A fan (located in the power supply) has been inserted.
LOG_ES_COOLING_ELEMENT_NORMAL	A fan (located in the power supply) has returned to normal status.
LOG_ES_COOLING_ELEMENT_REMOVED	A fan (located in the power supply) has been removed.
LOG_ES_COOLING_ELEMENT_WARNING	Enclosure cooling element has reported a warning condition.
LOG_ES_DISK_SLOT_ELEMENT_INSERTED	A device has been inserted into a disk slot.
LOG_ES_DISK_SLOT_ELEMENT_REMOVED	A device has been removed from a disk slot.
LOG_ES_ENCL_UPS_WARN_AC	UPS for the controller has started with no AC power.
LOG_ES_EXPANDER_ELEMENT_INSERTED	A DEM or an IO module has been inserted.
LOG_ES_EXPANDER_ELEMENT_NORMAL	The enclosure expander element has reported a warning condition.
LOG_ES_EXPANDER_ELEMENT_REMOVED	A DEM or an I/O module has been removed.
LOG_ES_EXPANDER_ELEMENT_WARNING	The DEM reports normal.
LOG_ES_POWER_SUPPLY_INSERTED	A power supply has been inserted.
LOG_ES_POWER_SUPPLY_NORMAL	The power supply has returned to a normal status.
LOG_ES_POWER_SUPPLY_REMOVED	A power supply has been removed.
LOG_ES_POWER_SUPPLY_WARNING	Power supply for the controller issued warning.
LOG_ES_TEMPERATURE_SENSOR_INSERTED	A temperature sensor (located in the power supply) has been inserted.
LOG_ES_TEMPERATURE_SENSOR_NORMAL	A temperature sensor (located in the power supply) has returned to normal status.
LOG_ES_TEMPERATURE_SENSOR_REMOVED	A temperature sensor (located in the power supply) has been removed.
LOG_ES_TEMPERATURE_SENSOR_WARNING	A temperature sensor in the enclosure has reported a warning condition.
LOG_ES_UPS_CLIENT_IS_RESPONSIVE	Controller enclosure UPS is responsive.
LOG_RAID_UNCORRECTED_MEDIUM_ERR	Aphysical device medium error could not be corrected due to lack of pool redundancy.
LOG_RAID_UNCORRECTED_SILENT_ERR	A SATAssure error (silent error) could not be corrected due to lack of pool redundancy.

Table 4. SNMP Only Events

SNMP ONLY Events	
Error	Description
LOG_RT_SNMP_TRAP_EVENT	Event received.
LOG_ST_ MEMBER_CHANGE	A member of a storage pool has changed its status.

Table 5. SNMP and Email Notification Events

SNMP and Email Notification Events	
Error	Description
LOG_AMPD_MPI_IOC_INIT_FAIL	SAS/SATA channel has failed to initialize.
LOG_ES_COOLING_ELEMENT_ERROR	Enclosure cooling element has reported an error condition.
LOG_ES_CTLR_PWR_SRC_CHANGED	Controller power source has changed.
LOG_ES_ENCL_UPS_WARN_BATT	UPS for the controller indicates that its battery should be replaced.
LOG_ES_ENCL_UPS_WARN_INTF	UPS for the controller indicates that its interface has failed. Check the cable.
LOG_ES_ENCL_UPS_WARN_UPS	UPS for the controller indicates that it has failed.
LOG_ES_EXPANDER_ELEMENT_ERROR	A DEM or an I/O module has reported an error condition.
LOG_ES_POWER_SUPPLY_ERROR	The power supply for one of the enclosure within the subsystem reported error condition.
LOG_ES_SET_BATTERY_CHARGER_FAILED	Attempt to set controller battery charger current and voltage failed.
LOG_ES_TEMPERATURE_SENSOR_ERROR	The enclosure temperature sensor has reported an error condition.
LOG_ES_UPS_CLIENT_IS_NOT_RESPONSIVE	Controller enclosure UPS is not responsive. Check the USB cable between the UPS and controller.
LOG_ST_HEARTBEAT_MISMATCH	Heartbeat mismatch reboot.
LOG_ST_MIR_STATE	Manual Intervention Required state.
LOG_ST_MIRROR_JOIN_FAIL	Other controller unable to flush mirror data status.
LOG_ST_NO_CONFIG_READ	Unable to read the configuration from backend drives status.
LOG_ST_NO_CONFIG_WRITE	Unable to write the configuration to backend drives status.
LOG_ST_POOL_CHANGE	Pool state changed.
LOG_ST_SET_AWL	Auto Write Lock condition.
LOG_ST_SET_CRITICAL	Critical condition state.
LOG_ST_SET_FAILED	Failed condition.
LOG_ST_SPLIT_BRAIN	The controllers have lost communications with each other and are operating in split-brain mode.
LOG_SYS_STARTUP	The system has been restarted.

Chapter 4

GUI Management Agent

This chapter provides information on using the SFA10000 GUI (Graphic User Interface) Management Agent.

4.1 Starting the GUI Management Agent

4.1.1 Login

Using a web browser, open a link to the IP address of the SFA10000 controller. The SFA Management System supports Mozilla FireFox (version 3.0.11 and above) as well as Microsoft Internet Explorer (version 8.0 and above).

At the login screen (Figure 112), enter the user name **admin** with the password **password** to log into the system. User name and passwords are case-sensitive.

Figure 112. GUI Login Screen

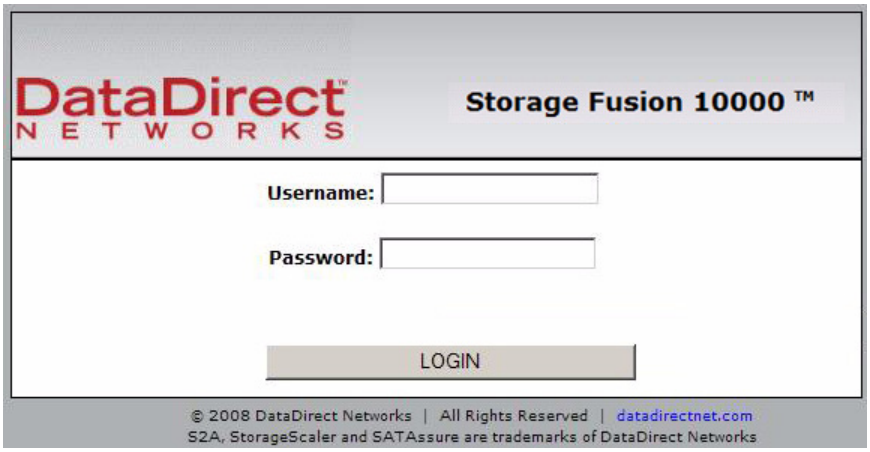

4.1.2 Home Screen and Health Indicators

The home screen shows the Health Indicator. Green color is an indication of healthy condition (Figure 113).

Figure 113. Home Screen (1)

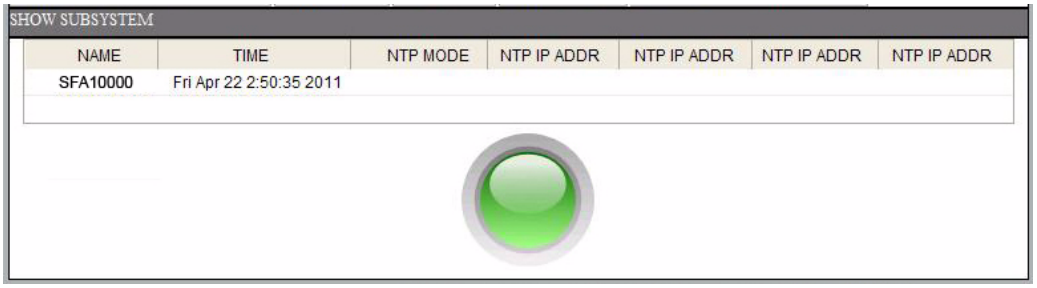

In Figure 114, the Health Indicator is orange indicating a faulty condition. The alert message below the Health Indicator indicates the component that requires attention.

Figure 114. Home Screen (2)

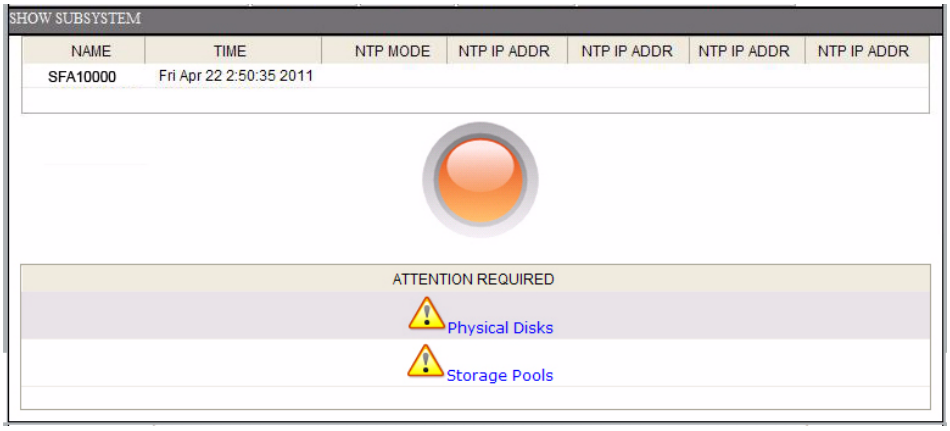

NOTE : Refer to Section 5.3, "Manual Intervention Required (MIR) States" on page 137 and the *SFA OS Service Manual* for specific instructions on how to clear the MIR states.

4.2 Subsystem Menu

Figure 115 shows the **Subsystem** menu.

Figure 115. Subsystem Menu

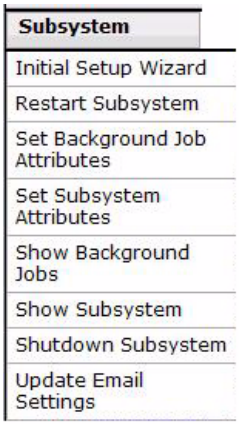

4.2.1 Show Subsystem

To display the current status of the subsystem, select **Subsystem > Show Subsystem**. Figure 116 illustration an example of the Show Subsystem screen.

It shows the subsystem's attributes. The color of the Health Indicator indicates the status of the subsystem. Green indicates a healthy condition. Orange indicates a faulty condition. The alert message below the orange Health Indicator indicates the component that requires attention.

Figure 116. Show Subsystem Screen

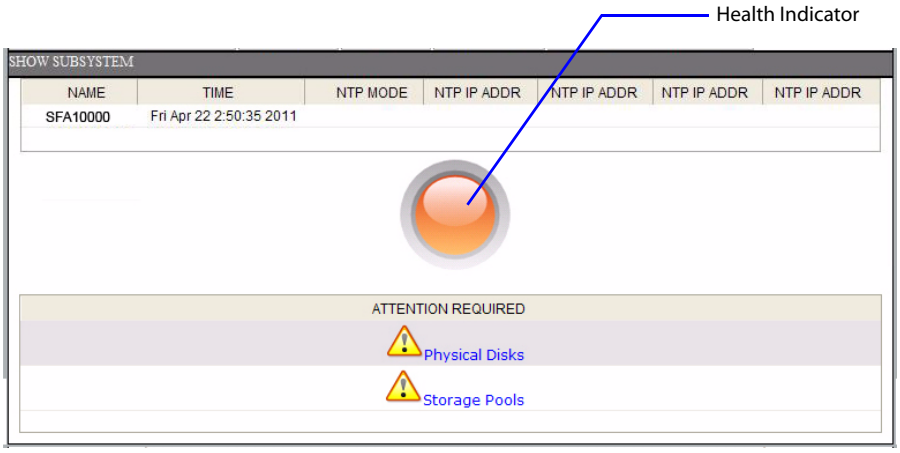

4.2.2 Set Subsystem Attributes

To change the subsystem's attributes, select **Subsystem > Set Subsystem Attributes**. Figure 117 illustrates the Set Subsystem Attributes screen. Make the desired changes and click **UPDATE** to save the changes.

Figure 117. Set Subsystem Attributes Screen

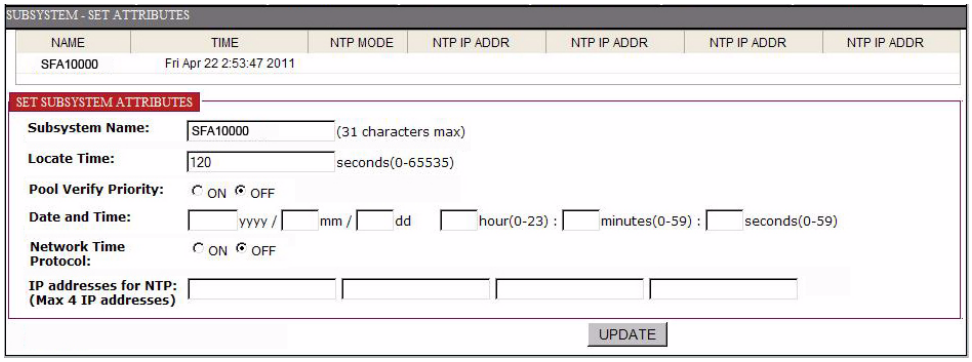

4.2.2.1 NTP Settings

NTP (Network Time Protocol) mode is available on the SFA10000. It provides a means for the controllers to synchronize their time across a network, usually within a small number of milliseconds over a long period of time. You can enter up to four NTP addresses as the time servers.

1. At the Set Subsystem Attributes screen (Figure 117), select ON to enable NTP.
2. In the **IP Addresses for NTP** field, enter the IP address of the time servers.
3. Click **UPDATE** to save the changes.

Please refer to Section 3.7, "Network Time Protocol Mode" on page 78 for information on behavioral changes when NTP mode is enabled.

4.2.3 Show Background Jobs

To monitor the progress of a job, select **Subsystem > Show Background Jobs**.

The Show Background Jobs screen (Figure 118) displays the list of current jobs. Click on the individual job name to display its detailed information.

Figure 118. Show Background Job Screen

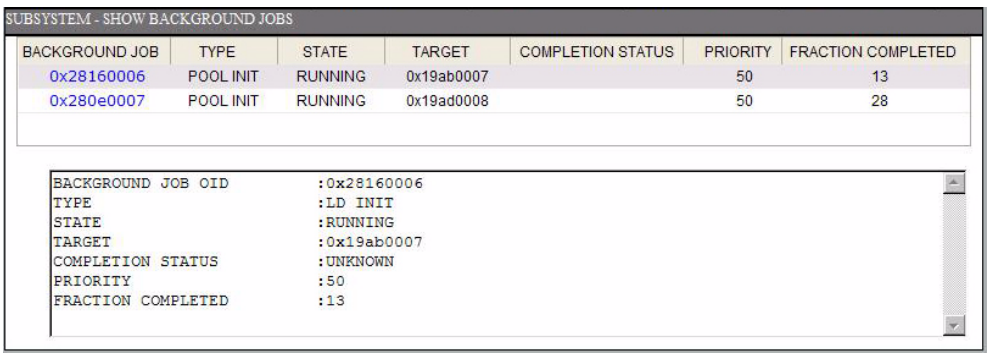

4.2.4 Set Background Jobs Attributes

The background job rates can be adjusted to minimize their impact on system performance. To set the job priority, select **Subsystem > Set Background Jobs**.

The Set Background Jobs screen (Figure 119) displays the list of current jobs. Select the job and the screen changes accordingly. Make the desired changes and click **UPDATE** to save.

Figure 119. Set Background Job Screen

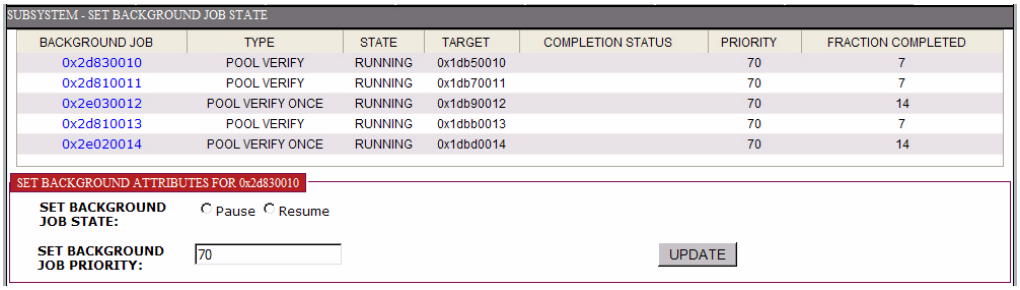

4.2.5 Email and Critical Event Notification Setup

Emails will be sent automatically as notification of a selected group of warning and error events that have occurred on the SFA10000. The recipient address is user-configurable and only one address is permitted per system.

To set up or change the email address:

Select **Subsystem > Update Email Settings**. The upper portion of the Update Email Settings screen (Figure 120) will display the current settings, if any.

Figure 120. Update Email Settings Screen

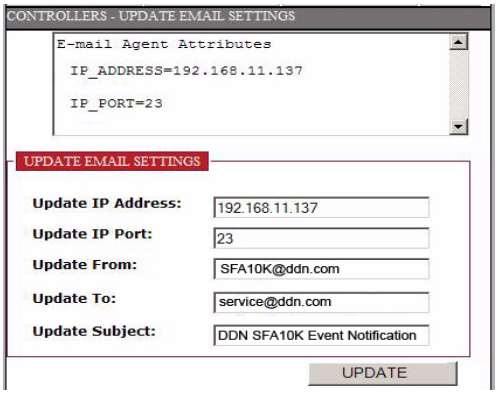

- In the **Update IP Address** field, enter the IP address of the SMTP server to which the SFA10000 will send email notifications.
- In the **Update IP Port** field, enter the IP port of the SMTP server to which the SFA10000 will send email notifications.

- In the **Update From** field (maximum field size 512), enter an identification string, such as the name of the controller. This is a string which will be included on the “From” line of the email notification message sent by the agent.
- In the **Update To** field (maximum field size 512), enter an Internet Mail Format (RFC2822) email address to which SFA10000 will send email notifications.
- In the **Update Subject** field (maximum field size 512), enter a string to be included on the subject line of the email notification sent by the email agent.

Click **UPDATE** to save the changes.

The tables listed in Section 3.13.5.3, "Inquiry Items and Events" on page 103 detail the additional inquiry items and events that will be trapped in the SNMP MIB and in the Email Agent.

4.3 Controllers Menu

Figure 121 shows the **Controllers** pull down menu.

Figure 121. Controllers Menu

4.3.1 Show Controllers

The **Show Controllers** function display the list of controllers installed in the system (Figure 122). If a Manual Intervention Required condition is present, click on the highlighted event to get detailed information.

Figure 122. Show All Controllers Screen

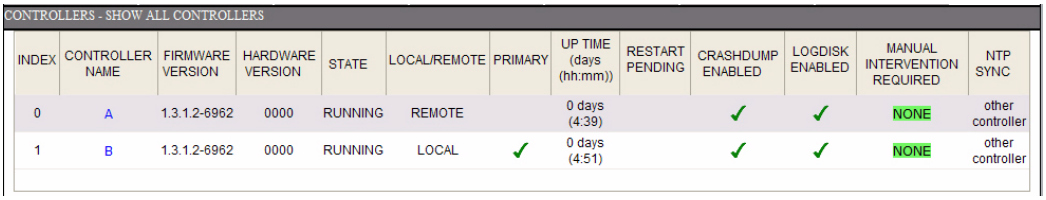

4.3.2 Set Controller Attributes

To change the controller's attributes, select **Controllers > Set Controller Attributes**. Figure 123 illustrates the Set Controller Attributes screen. Make the desired changes and click **UPDATE** to save the changes.

Figure 123. Set Controller Attributes Screen

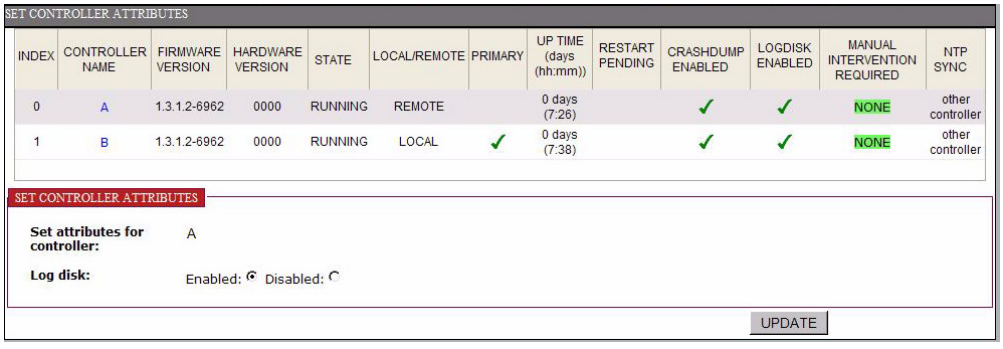

4.3.3 Update Controller Firmware

You can update the controller firmware using the **Update Firmware** function. After you have obtained the new firmware files, select **Controllers > Update Firmware**. At the Update Firmware screen (Figure 124), locate the firmware file and click **UPDATE FIRMWARE**. Then follow the onscreen instructions to complete the update. Also refer to Section 3.11.2, "Controller Firmware Update Procedure" on page 89 for detailed instructions.

Figure 124. Update Controller Firmware Screen

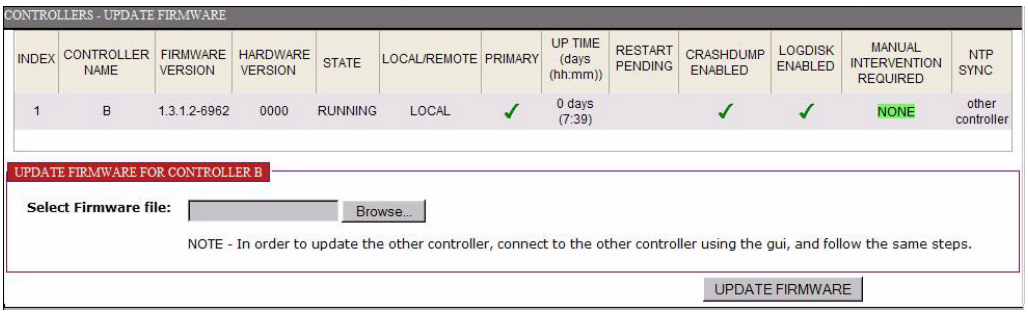

4.3.4 View System Logs

To view the system logs, select **Controllers > View Log**. At the View Log Entries screen (Figure 125), select the display options and click **VIEW LOGS** to display the log (Figure 126).

Figure 125. View Log Entries Screen

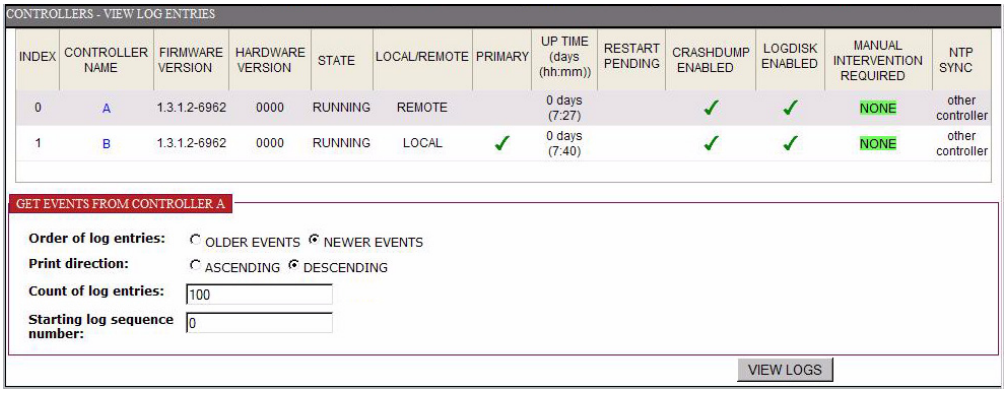

Figure 126. View Log Screen

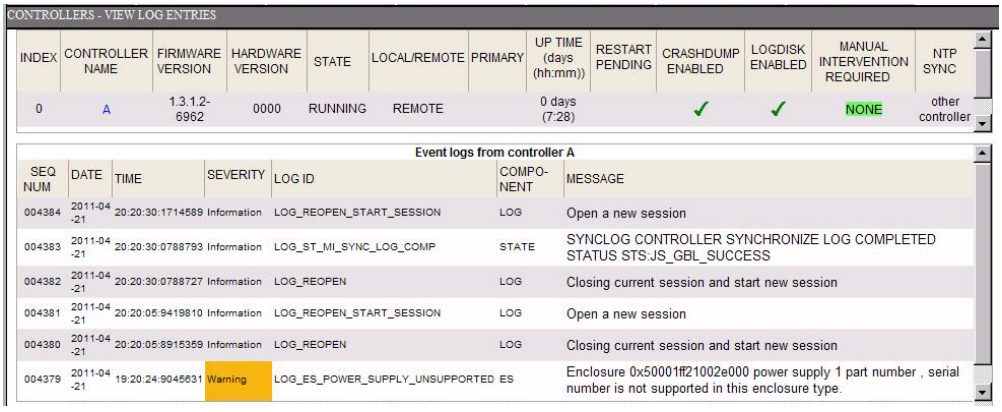

4.4 Physical Disks Menu

Figure 127 shows the **Physical Disks** pull down menu.

Figure 127. Physical Disks Menu

4.4.1 Show List of Physical Disks

The **Show All Physical Disks** function displays the list of disks that are installed in the system (Figure 128). Click on the disk to display its detailed information.

NOTE : If the disks are not visible when the system first starts up, select **Subsystem Restart** from the Subsystem Menu to reboot the system.

Figure 128. Show Physical Disks Screen

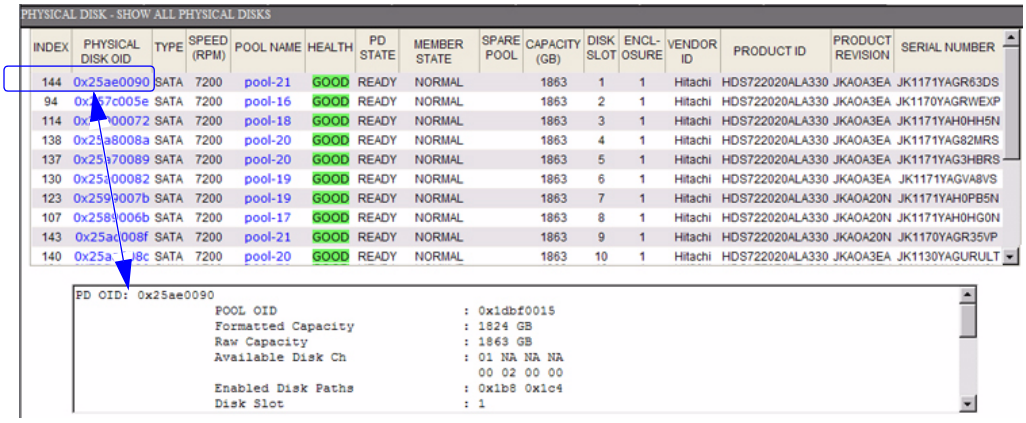
PHYSICAL DISK - SHOW ALL PHYSICAL DISKS
INDEX PHYSICAL DISK OID TYPE SPEED (RPM) POOL NAME HEALTH PD STATE MEMBER STATE SPARE POOL CAPACITY (GB) DISK SLOT ENCL-OSURE VENDOR ID PRODUCT ID PRODUCT REVISION SERIAL NUMBER
144 0x25ae0090 SATA 7200 pool-21 GOOD READY NORMAL 1863 1 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAGR63DS
94 0x25c005e SATA 7200 pool-16 GOOD READY NORMAL 1863 2 1 Hitachi HDS722020ALA330 JKAOA3EA JK1170YAGRWEXP
114 0x 00072 SATA 7200 pool-18 GOOD READY NORMAL 1863 3 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAH0HH5N
138 0x25a8008a SATA 7200 pool-20 GOOD READY NORMAL 1863 4 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAG82MRS
137 0x25a70089 SATA 7200 pool-20 GOOD READY NORMAL 1863 5 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAG3HBRS
130 0x25a00082 SATA 7200 pool-19 GOOD READY NORMAL 1863 6 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAGVA8VS
123 0x2590007b SATA 7200 pool-19 GOOD READY NORMAL 1863 7 1 Hitachi HDS722020ALA330 JKAOA20N JK1171YAH0PB5N
107 0x2580006b SATA 7200 pool-17 GOOD READY NORMAL 1863 8 1 Hitachi HDS722020ALA330 JKAOA20N JK1171YAH0HG0N
143 0x25ad008f SATA 7200 pool-21 GOOD READY NORMAL 1863 9 1 Hitachi HDS722020ALA330 JKAOA20N JK1170YAGR35VP
140 0x25a 8c SATA 7200 pool-20 GOOD READY NORMAL 1863 10 1 Hitachi HDS722020ALA330 JKAOA3EA JK1130YAGURULT
PD OID: 0x25ae0090
POOL OID : 0x1dbf0015
Formatted Capacity : 1824 GB
Raw Capacity : 1863 GB
Available Disk Ch : 01 NA NA NA
00 02 00 00
Enabled Disk Paths : 0x1b8 0x1c4
Disk Slot : 1

NOTE : If asterisks appear in the Physical Disk column when you attempt to view the advanced disk information or attempt to locate a disk, this shows that only one controller sees the disk. Therefore, a controller is down or another hardware issue exists somewhere.

4.4.2 Locate Disks

To locate a disk in the enclosure, select **Physical Disk > Locate Disks**. In the Locate screen (Figure 129), click to select the disk. The lower half of the Locate screen displays the disk slot into which the disk is installed. Click **LOCATE** to flash the status LED of the disk module.

Figure 129. Locate Disks Screen

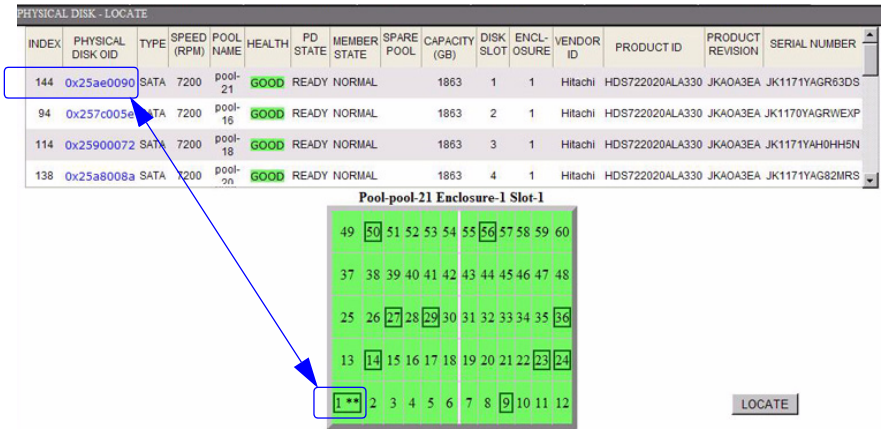
PHYSICAL DISK - LOCATE
INDEX
PHYSICAL DISK OID
TYPE
SPEED (RPM)
POOL NAME
HEALTH
PD STATE
MEMBER STATE
SPARE POOL
CAPACITY (GB)
DISK SLOT
ENCL-OSURE
VENDOR ID
PRODUCT ID
PRODUCT REVISION
SERIAL NUMBER
144 0x25ae0090 SATA 7200 pool-21 GOOD READY NORMAL 1863 1 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAGR63DS
94 0x257c005e SATA 7200 pool-16 GOOD READY NORMAL 1863 2 1 Hitachi HDS722020ALA330 JKAOA3EA JK1170YAGRWEXP
114 0x25900072 SATA 7200 pool-18 GOOD READY NORMAL 1863 3 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAH0HH5N
138 0x25a8008a SATA 7200 pool-20 GOOD READY NORMAL 1863 4 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAG82MRS
Pool-pool-21 Enclosure-1 Slot-1
49 50 51 52 53 54 55 56 57 58 59 60
37 38 39 40 41 42 43 44 45 46 47 48
25 26 27 28 29 30 31 32 33 34 35 36
13 14 15 16 17 18 19 20 21 22 23 24
1 ** 2 3 4 5 6 7 8 9 10 11 12
LOCATE

4.4.3 Remove Disk from Spare Pool

If you need to remove a disk from a spare pool, select **Physical Disks > Remove from Spare**. (Figure 130). Select the disk and click **REMOVE**.

Figure 130. Remove Disks Screen

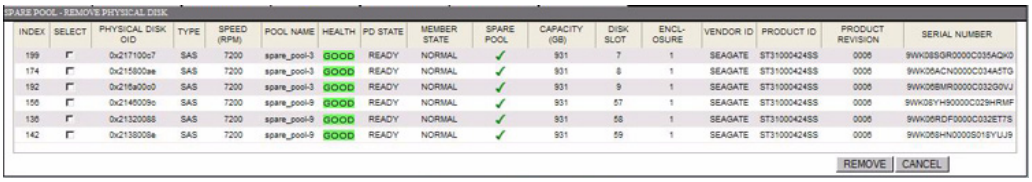
SPARE POOL - REMOVE PHYSICAL DISK
INDEX
SELECT
PHYSICAL DISK OID
TYPE
SPEED (RPM)
POOL NAME
HEALTH
PD STATE
MEMBER STATE
SPARE POOL
CAPACITY (GB)
DISK SLOT
ENCL-OSURE
VENDOR ID
PRODUCT ID
PRODUCT REVISION
SERIAL NUMBER
199 0x217100c7 SAS 7200 spare_pool-3 GOOD READY NORMAL 931 7 1 SEAGATE ST3100424SS 0006 9WK08SGR0000C035AQK0
174 0x215800ae SAS 7200 spare_pool-3 GOOD READY NORMAL 931 8 1 SEAGATE ST3100424SS 0006 9WK05ACN0000C034A5TG
192 0x216a00c0 SAS 7200 spare_pool-3 GOOD READY NORMAL 931 9 1 SEAGATE ST3100424SS 0006 9WK05BMR0000C032G0VJ
156 0x2146009c SAS 7200 spare_pool-9 GOOD READY NORMAL 931 57 1 SEAGATE ST3100424SS 0006 9WK0EYH90000C029HRMF
136 0x21320088 SAS 7200 spare_pool-9 GOOD READY NORMAL 931 58 1 SEAGATE ST3100424SS 0006 9WK06RDF0000C032ET7S
142 0x2138008e SAS 7200 spare_pool-9 GOOD READY NORMAL 931 59 1 SEAGATE ST3100424SS 0006 9WK06SHN0000S018YUJ9
REMOVE
CANCEL

4.4.4 Assign Disk to Pools

You can add a disk to a spare pool or a storage pool at any time, select **Physical Disks > Assign to Pool**. (Figure 131). Select the disk(s) and the pool. Click **ASSIGN TO POOL** to save the changes.

Figure 131. Assign Disks to Pool Screen

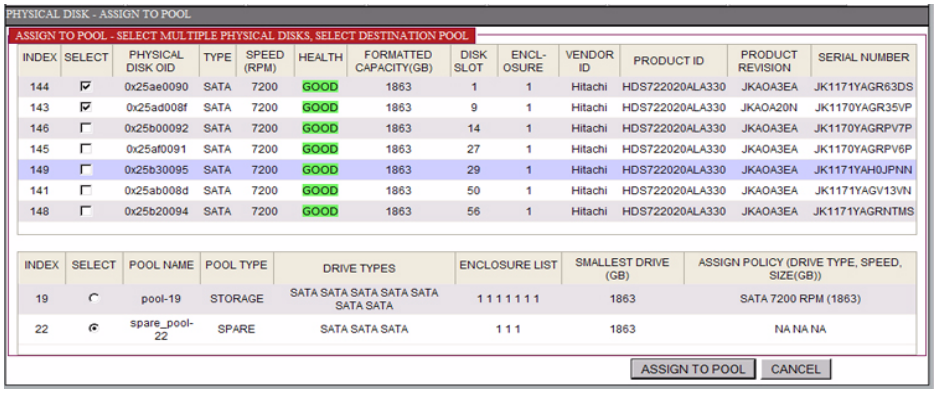
PHYSICAL DISK - ASSIGN TO POOL
ASSIGN TO POOL - SELECT MULTIPLE PHYSICAL DISKS, SELECT DESTINATION POOL
INDEX
SELECT
PHYSICAL DISK OID
TYPE
SPEED (RPM)
HEALTH
FORMATTED CAPACITY(GB)
DISK SLOT
ENCL-OSURE
VENDOR ID
PRODUCT ID
PRODUCT REVISION
SERIAL NUMBER
144 0x25ae0090 SATA 7200 GOOD 1863 1 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAGR63DS
143 0x25ad008f SATA 7200 GOOD 1863 9 1 Hitachi HDS722020ALA330 JKAOA20N JK1170YAGR35VP
146 0x25b00092 SATA 7200 GOOD 1863 14 1 Hitachi HDS722020ALA330 JKAOA3EA JK1170YAGRPV7P
145 0x25af0091 SATA 7200 GOOD 1863 27 1 Hitachi HDS722020ALA330 JKAOA3EA JK1170YAGRPV6P
149 0x25b30095 SATA 7200 GOOD 1863 29 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAH0JPNN
141 0x25ab008d SATA 7200 GOOD 1863 50 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAGV13VN
148 0x25b20094 SATA 7200 GOOD 1863 56 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAGRNTMS
INDEX
SELECT
POOL NAME
POOL TYPE
DRIVE TYPES
ENCLOSURE LIST
SMALLEST DRIVE (GB)
ASSIGN POLICY (DRIVE TYPE, SPEED, SIZE(GB))
19 pool-19 STORAGE SATA SATA SATA SATA SATA SATA SATA 1 1 1 1 1 1 1 1863 SATA 7200 RPM (1863)
22 spare_pool-22 SPARE SATA SATA SATA 1 1 1 1863 NA NA NA
ASSIGN TO POOL
CANCEL

4.5 Unassigned Pool

By default, all the disks or replacement disks are initially allocated to this pool. When a disk has failed or been released (from a spare pool, for example), it is also moved to this pool.

4.5.1 Show Unassigned Pool

Select **Pools > Show Unassigned Pool** (Figure 132).

Figure 132. Pools Menu

At the Show Unassigned Pool screen (Figure 133), click on **Total Members** to display the list of disks in this pool.

Figure 133. Show Unassigned Pool Screen

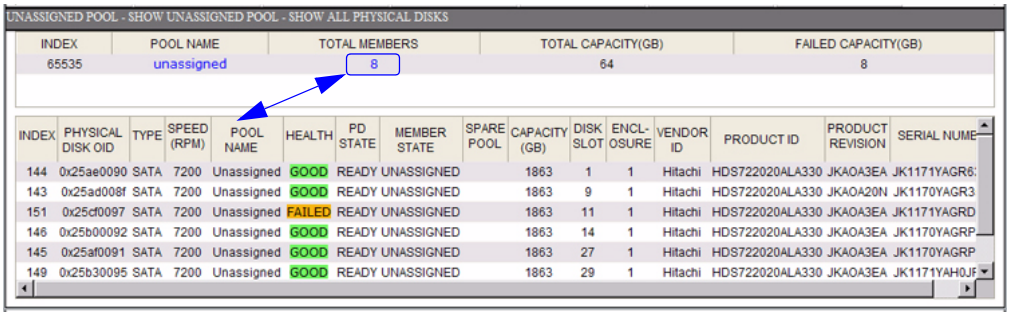
UNASSIGNED POOL - SHOW UNASSIGNED POOL - SHOW ALL PHYSICAL DISKS
INDEX
POOL NAME
TOTAL MEMBERS
TOTAL CAPACITY(GB)
FAILED CAPACITY(GB)
65535
unassigned
8
64
8
INDEX PHYSICAL DISK OID TYPE SPEED (RPM) POOL NAME HEALTH PD STATE MEMBER STATE SPARE POOL CAPACITY (GB) DISK SLOT ENCL-OSURE VENDOR ID PRODUCT ID PRODUCT REVISION SERIAL NUME
144 0x25ae0090 SATA 7200 Unassigned GOOD READY UNASSIGNED 1863 1 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAGR6
143 0x25ad008f SATA 7200 Unassigned GOOD READY UNASSIGNED 1863 9 1 Hitachi HDS722020ALA330 JKAOA20N JK1170YAGR3
151 0x25cf0097 SATA 7200 Unassigned FAILED READY UNASSIGNED 1863 11 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAGRD
146 0x25b00092 SATA 7200 Unassigned GOOD READY UNASSIGNED 1863 14 1 Hitachi HDS722020ALA330 JKAOA3EA JK1170YAGRP
145 0x25af0091 SATA 7200 Unassigned GOOD READY UNASSIGNED 1863 27 1 Hitachi HDS722020ALA330 JKAOA3EA JK1170YAGRP
149 0x25b30095 SATA 7200 Unassigned GOOD READY UNASSIGNED 1863 29 1 Hitachi HDS722020ALA330 JKAOA3EA JK1171YAH0JF

4.5.2 Locate Unassigned Pool

To locate the disks in the unassigned pool, select **Pools > Locate Unassigned Pool**.

Click “unassigned” to display the disk map (Figure 134).

- Green color indicates a healthy disk that is available.
- White color indicates a healthy disk that has been assigned to either a storage pool or spare pool.
- Orange color indicates that the disk has failed.

Click **LOCATE** to flash the status LED on the unassigned disks.

Figure 134. Locate Unassigned Pool Screen

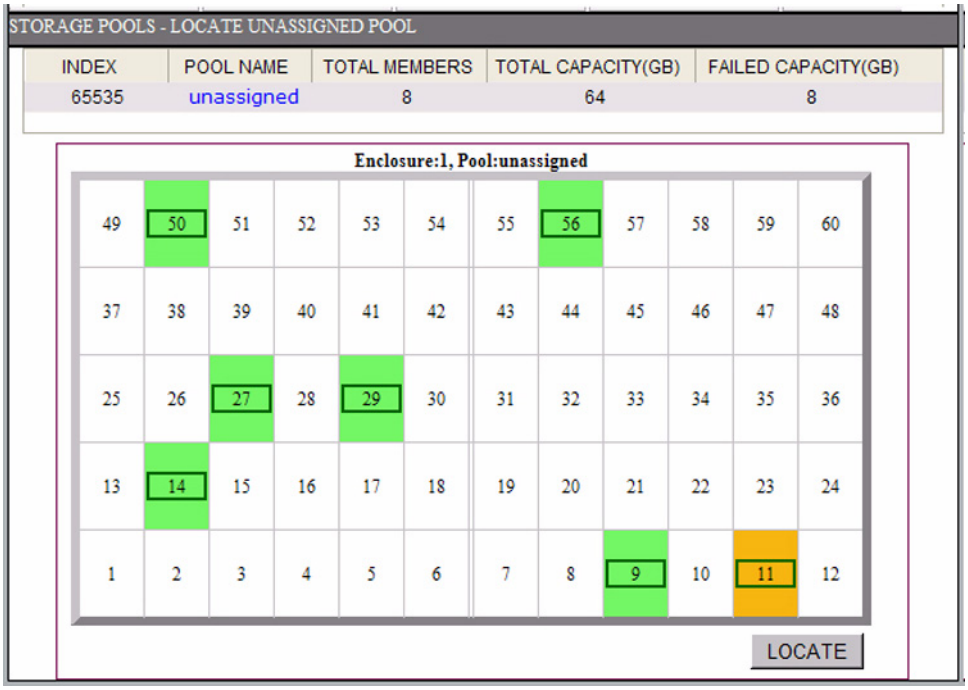
STORAGE POOLS - LOCATE UNASSIGNED POOL
INDEX
POOL NAME
TOTAL MEMBERS
TOTAL CAPACITY(GB)
FAILED CAPACITY(GB)
65535
unassigned
8
64
8
Enclosure:1, Pool:unassigned
49 50 51 52 53 54 55 56 57 58 59 60
37 38 39 40 41 42 43 44 45 46 47 48
25 26 27 28 29 30 31 32 33 34 35 36
13 14 15 16 17 18 19 20 21 22 23 24
1 2 3 4 5 6 7 8 9 10 11 12
LOCATE

4.6 Storage Pools

Storage pool commands are located under the **Pools** pull down menu (Figure 135):

Figure 135. Pools Menu

4.6.1 Show List of Storage Pools

To display the list of existing storage pools, select **Pools > Show Storage Pools**. The list of storage pools is displayed (Figure 136). Click on the individual storage pool name to display its detailed information.

Figure 136. Storage Pool List

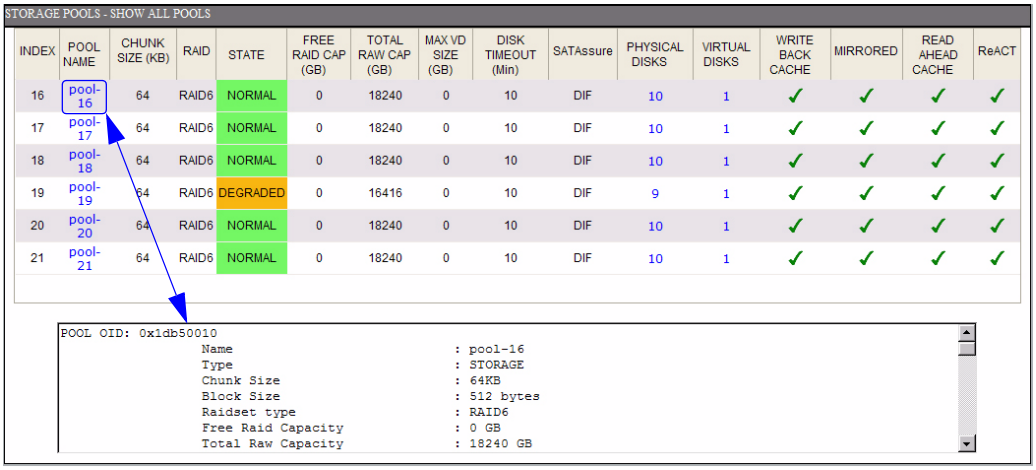
STORAGE POOLS - SHOW ALL POOLS
INDEX
POOL NAME
CHUNK SIZE (KB)
RAID
STATE
FREE RAID CAP (GB)
TOTAL RAW CAP (GB)
MAX VD SIZE (GB)
DISK TIMEOUT (Min)
SATAssure
PHYSICAL DISKS
VIRTUAL DISKS
WRITE BACK CACHE
MIRRORED
READ AHEAD CACHE
ReACT
16 pool-16 64 RAID6 NORMAL 0 18240 0 10 DIF 10 1 ✓ ✓ ✓ ✓
17 pool-17 64 RAID6 NORMAL 0 18240 0 10 DIF 10 1 ✓ ✓ ✓ ✓
18 pool-18 64 RAID6 NORMAL 0 18240 0 10 DIF 10 1 ✓ ✓ ✓ ✓
19 pool-19 64 RAID6 DEGRADED 0 16416 0 10 DIF 9 1 ✓ ✓ ✓ ✓
20 pool-20 64 RAID6 NORMAL 0 18240 0 10 DIF 10 1 ✓ ✓ ✓ ✓
21 pool-21 64 RAID6 NORMAL 0 18240 0 10 DIF 10 1 ✓ ✓ ✓ ✓
POOL OID: 0x1db50010
Name : pool-16
Type : STORAGE
Chunk Size : 64KB
Block Size : 512 bytes
Raidset type : RAID6
Free Raid Capacity : 0 GB
Total Raw Capacity : 18240 GB

4.6.2 Create a Storage Pool

A storage pool on a SFA10000 has the following attributes:

- RAID Level
  Storage pools can be configured to use either a RAID 1, RAID 5 or RAID 6 parity scheme. In RAID 1, the capacity of one disk is used for data duplication. In RAID 5, the capacity of one disk is reserved for parity, allowing data recovery in the event of a single disk loss in the storage pool. In RAID 6, the capacity of two disks is reserved for parity, allowing data recovery if either one or two disks are down in a storage pool. ***For maximum data protection, DataDirect Networks recommends the use of RAID 6***.

- Chunk Size
  The chunk size (in KiB blocks) defines the amount of data written to a single disk before proceeding to the next disk in the storage pool.

> **NOTE :** RAID1 is a two member RAID set where the data is mirrored on each disk. There is no parity, hence, the chunk size is fixed.

- Disk Count
  A RAID1 storage pool may consists of 2 physical disks. A RAID5 storage pool may consist of 5 or 9 physical disks. A RAID6 storage pool may consist of 6 or 10 physical disks. For maximum performance, select disks with the same characteristics (such as SAS/SATA, capacity, and RPM).

- Drive Type
  Drive types can be SAS or SATA.

- Drive Size
  Drive size is the capacity of the disk.

- Spindle Speed
  You may choose (RPM): 15000, 10000, 7200, 5400, 0 or SSD.

- SATAssure
  SATAssure technology is designed by DDN to improve the reliability of enterprise SATA disks and make sure that data integrity is always mentioned for all I/O operations.

To create a storage pool, select **Pools > Create Storage Pool**. At the Create Storage Pool screen (Figure 137), select all the attributes for the pool. Click **CREATE STORAGE POOL** to create the pool.

Figure 137. Create Storage Pool Screen

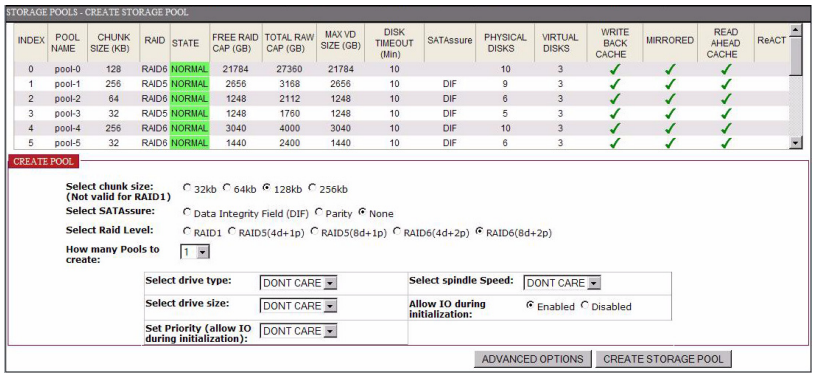
STORAGE POOLS - CREATE STORAGE POOL
INDEX
POOL NAME
CHUNK SIZE (KB)
RAID
STATE
FREE RAID CAP (GB)
TOTAL RAW CAP (GB)
MAX VD SIZE (GB)
DISK TIMEOUT (Min)
SATAssure
PHYSICAL DISKS
VIRTUAL DISKS
WRITE BACK CACHE
MIRRORED
READ AHEAD CACHE
ReACT
0 pool-0 128 RAID6 NORMAL 21784 27360 21784 10 10 3
1 pool-1 256 RAID5 NORMAL 2656 3168 2656 10 DIF 9 3
2 pool-2 64 RAID6 NORMAL 1248 2112 1248 10 DIF 6 3
3 pool-3 32 RAID5 NORMAL 1248 1760 1248 10 DIF 5 3
4 pool-4 256 RAID6 NORMAL 3040 4000 3040 10 DIF 10 3
5 pool-5 32 RAID6 NORMAL 1440 2400 1440 10 DIF 6 3
CREATE POOL
Select chunk size: (Not valid for RAID1)
32kb 64kb 128kb 256kb
Select SATAssure:
Data Integrity Field (DIF) Parity None
Select Raid Level:
RAID1 RAID5(4d+1p) RAID5(8d+1p) RAID6(4d+2p) RAID6(8d+2p)
How many Pools to create:
Select drive type:
DONT CARE
Select spindle Speed:
DONT CARE
Select drive size:
DONT CARE
Allow IO during initialization:
Enabled Disabled
Set Priority (allow IO during initialization):
DONT CARE
ADVANCED OPTIONS
CREATE STORAGE POOL

OPTIONAL: You may explicitly select the disks to add to the pool using the **ADVANCED OPTIONS**. At the Create Storage Pool Advanced screen (Figure 138), select the disks to add to the pool to be created. Otherwise, click **CANCEL** to cancel or click **BACK** to return to the previous screen.

Figure 138. Select Disks for Storage Pool

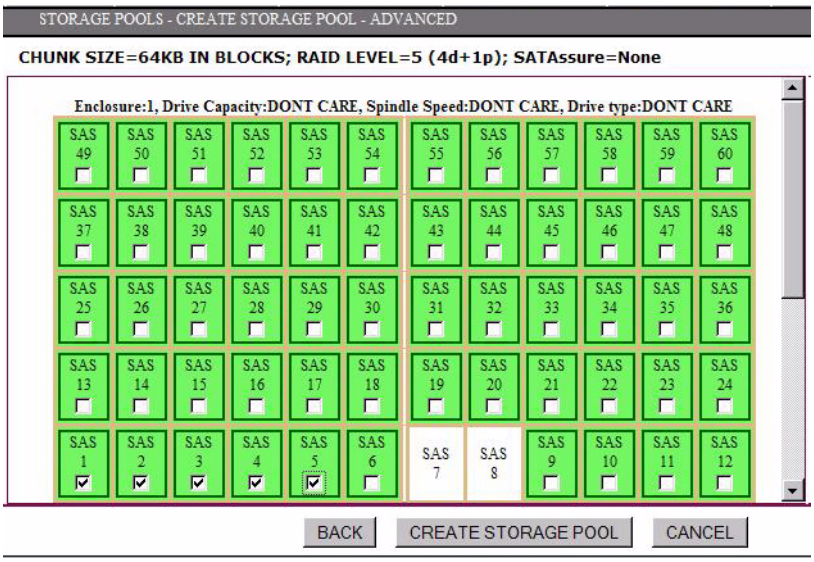
STORAGE POOLS - CREATE STORAGE POOL - ADVANCED
CHUNK SIZE=64KB IN BLOCKS; RAID LEVEL=5 (4d+1p); SATAssure=None
Enclosure:1, Drive Capacity:DONT CARE, Spindle Speed:DONT CARE, Drive type:DONT CARE
SAS 49 SAS 50 SAS 51 SAS 52 SAS 53 SAS 54 SAS 55 SAS 56 SAS 57 SAS 58 SAS 59 SAS 60
SAS 37 SAS 38 SAS 39 SAS 40 SAS 41 SAS 42 SAS 43 SAS 44 SAS 45 SAS 46 SAS 47 SAS 48
SAS 25 SAS 26 SAS 27 SAS 28 SAS 29 SAS 30 SAS 31 SAS 32 SAS 33 SAS 34 SAS 35 SAS 36
SAS 13 SAS 14 SAS 15 SAS 16 SAS 17 SAS 18 SAS 19 SAS 20 SAS 21 SAS 22 SAS 23 SAS 24
SAS 1 SAS 2 SAS 3 SAS 4 SAS 5 SAS 6 SAS 7 SAS 8 SAS 9 SAS 10 SAS 11 SAS 12
BACK
CREATE STORAGE POOL
CANCEL

4.6.3 Configure Storage Pool Attributes

To change the attributes of a storage pool, select **Pools > Set Storage Pool Attributes**.

At the Set Attributes for Pool screen (Figure 139), make the desired changes and click **UPDATE** to save them.

Figure 139. Set Storage Pool Attributes Screen

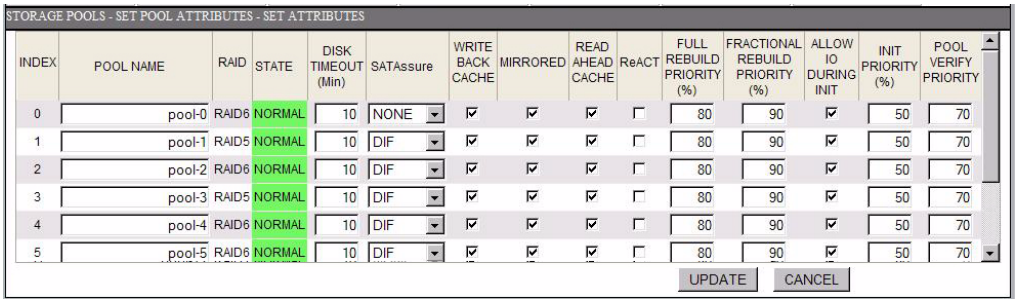
STORAGE POOLS - SET POOL ATTRIBUTES - SET ATTRIBUTES
INDEX POOL NAME RAID STATE DISK TIMEOUT (Min) SATAssure WRITE BACK CACHE MIRRORED READ AHEAD CACHE ReACT FULL REBUILD PRIORITY (%) FRACTIONAL REBUILD PRIORITY (%) ALLOW IO DURING INIT INIT PRIORITY (%) POOL VERIFY PRIORITY
0 pool-0 RAID6 NORMAL 10 NONE 80 90 50 70
1 pool-1 RAID5 NORMAL 10 DIF 80 90 50 70
2 pool-2 RAID6 NORMAL 10 DIF 80 90 50 70
3 pool-3 RAID5 NORMAL 10 DIF 80 90 50 70
4 pool-4 RAID6 NORMAL 10 DIF 80 90 50 70
5 pool-5 RAID6 NORMAL 10 DIF 80 90 50 70
UPDATE
CANCEL

4.6.4 Locate Storage Pools

To locate the disks in a storage pool, select **Pools > Locate Storage Pool**.

At the Locate Pool screen (Figure 140), select the storage pool that you want to locate. The screen changes accordingly to show the pool's disks. Click **LOCATE** to flash the status LED of those disks.

Figure 140. Locate Storage Pool Screen

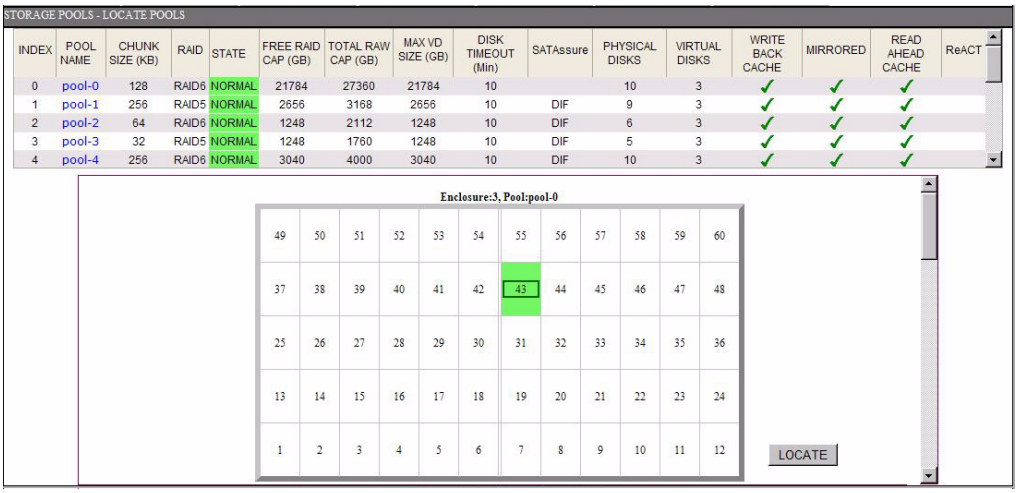
STORAGE POOLS - LOCATE POOLS
INDEX
POOL NAME
CHUNK SIZE (KB)
RAID
STATE
FREE RAID CAP (GB)
TOTAL RAW CAP (GB)
MAX VD SIZE (GB)
DISK TIMEOUT (Min)
SATAssure
PHYSICAL DISKS
VIRTUAL DISKS
WRITE BACK CACHE
MIRRORED
READ AHEAD CACHE
ReACT
0 pool-0 128 RAID6 NORMAL 21784 27360 21784 10 10 3
1 pool-1 256 RAID5 NORMAL 2656 3168 2656 10 DIF 9 3
2 pool-2 64 RAID6 NORMAL 1248 2112 1248 10 DIF 6 3
3 pool-3 32 RAID5 NORMAL 1248 1760 1248 10 DIF 5 3
4 pool-4 256 RAID6 NORMAL 3040 4000 3040 10 DIF 10 3
Enclosure:3, Pool:pool-0
49 50 51 52 53 54 55 56 57 58 59 60
37 38 39 40 41 42 43 44 45 46 47 48
25 26 27 28 29 30 31 32 33 34 35 36
13 14 15 16 17 18 19 20 21 22 23 24
1 2 3 4 5 6 7 8 9 10 11 12
LOCATE

4.6.5 Delete a Storage Pool

If you need to delete a storage pool, select **Pools > Delete Storage Pool**.

Select the pool that you want to delete (Figure 141). Click **DELETE** to delete that storage pool. You will be prompt to confirm deletion, click **OK** to proceed or **CANCEL** to cancel operation.

Figure 141. Delete Storage Pool Screen

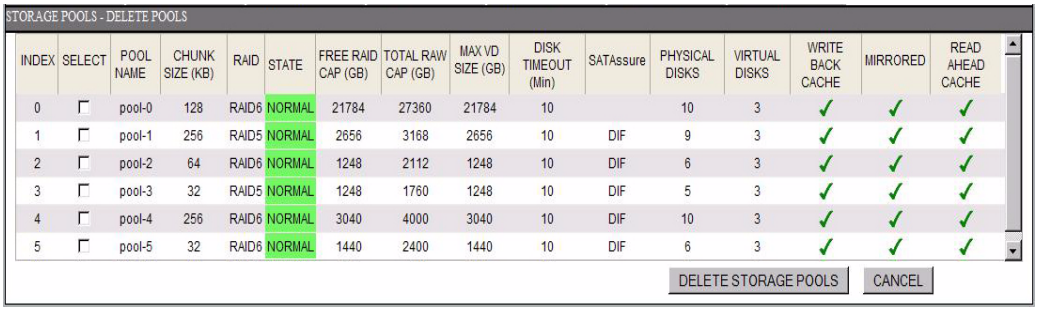
STORAGE POOLS - DELETE POOLS
INDEX
SELECT
POOL NAME
CHUNK SIZE (KB)
RAID
STATE
FREE RAID CAP (GB)
TOTAL RAW CAP (GB)
MAX VD SIZE (GB)
DISK TIMEOUT (Min)
SATAssure
PHYSICAL DISKS
VIRTUAL DISKS
WRITE BACK CACHE
MIRRORED
READ AHEAD CACHE
0 pool-0 128 RAID6 NORMAL 21784 27360 21784 10 10 3
1 pool-1 256 RAID5 NORMAL 2656 3168 2656 10 DIF 9 3
2 pool-2 64 RAID6 NORMAL 1248 2112 1248 10 DIF 6 3
3 pool-3 32 RAID5 NORMAL 1248 1760 1248 10 DIF 5 3
4 pool-4 256 RAID6 NORMAL 3040 4000 3040 10 DIF 10 3
5 pool-5 32 RAID6 NORMAL 1440 2400 1440 10 DIF 6 3
DELETE STORAGE POOLS
CANCEL

4.7 Virtual Disks Menu

A Virtual Disk is the storage unit presented to any attached host. A virtual disk can be created to use all or just a part of the capacity of a single storage pool. VDs are created in increments of 8 GiB. For example, 16 GiB of storage space will be allocated when creating a VD of 10 GiB.

The virtual disk commands are located under the **Virtual Disk** pull down menu (Figure 142):

Figure 142. Virtual Disk Menu

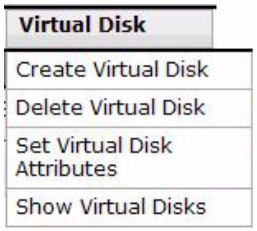

4.7.1 Show List of Virtual Disks

Select **Show Virtual Disks** from the **Virtual Disk** Menu (Figure 142). The list of configured virtual disk(s) appears (Figure 143). Click on the individual VD name to display its detailed information.

Figure 143. Show Virtual Disk Details Screen

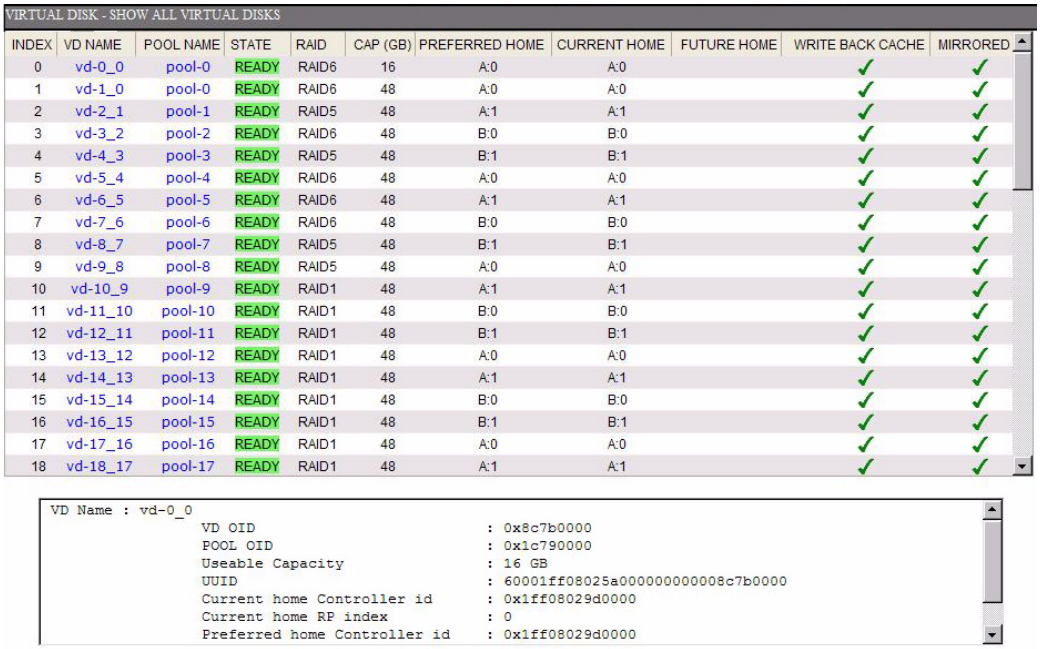
VIRTUAL DISK - SHOW ALL VIRTUAL DISKS
INDEX VD NAME POOL NAME STATE RAID CAP (GB) PREFERRED HOME CURRENT HOME FUTURE HOME WRITE BACK CACHE MIRRORED
0 vd-0_0 pool-0 READY RAID6 16 A:0 A:0
1 vd-1_0 pool-0 READY RAID6 48 A:0 A:0
2 vd-2_1 pool-1 READY RAID5 48 A:1 A:1
3 vd-3_2 pool-2 READY RAID6 48 B:0 B:0
4 vd-4_3 pool-3 READY RAID5 48 B:1 B:1
5 vd-5_4 pool-4 READY RAID6 48 A:0 A:0
6 vd-6_5 pool-5 READY RAID6 48 A:1 A:1
7 vd-7_6 pool-6 READY RAID6 48 B:0 B:0
8 vd-8_7 pool-7 READY RAID5 48 B:1 B:1
9 vd-9_8 pool-8 READY RAID5 48 A:0 A:0
10 vd-10_9 pool-9 READY RAID1 48 A:1 A:1
11 vd-11_10 pool-10 READY RAID1 48 B:0 B:0
12 vd-12_11 pool-11 READY RAID1 48 B:1 B:1
13 vd-13_12 pool-12 READY RAID1 48 A:0 A:0
14 vd-14_13 pool-13 READY RAID1 48 A:1 A:1
15 vd-15_14 pool-14 READY RAID1 48 B:0 B:0
16 vd-16_15 pool-15 READY RAID1 48 B:1 B:1
17 vd-17_16 pool-16 READY RAID1 48 A:0 A:0
18 vd-18_17 pool-17 READY RAID1 48 A:1 A:1
VD Name : vd-0_0
VD OID : 0x8c7b0000
POOL OID : 0x1c790000
Useable Capacity : 16 GB
UUID : 60001ff08025a000000000008c7b0000
Current home Controller id : 0x1ff08029d0000
Current home RP index : 0
Preferred home Controller id : 0x1ff08029d0000

4.7.2 Create a Virtual Disk

To create a virtual disk, select **Virtual Disks > Create Virtual Disk**.

At the Create Virtual Disk screen (Figure 144), select the storage pool(s) to be used. Then enter the capacity in GB. Click **CREATE VIRTUAL DISK** to create the virtual disk.

Figure 144. Create Virtual Disk Screen

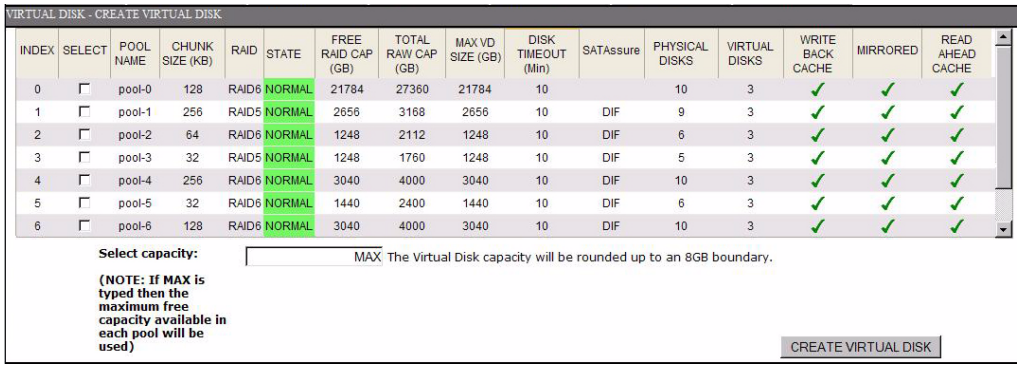
VIRTUAL DISK - CREATE VIRTUAL DISK
INDEX
SELECT
POOL NAME
CHUNK SIZE (KB)
RAID
STATE
FREE RAID CAP (GB)
TOTAL RAW CAP (GB)
MAX VD SIZE (GB)
DISK TIMEOUT (Min)
SATAssure
PHYSICAL DISKS
VIRTUAL DISKS
WRITE BACK CACHE
MIRRORED
READ AHEAD CACHE
0 pool-0 128 RAID6 NORMAL 21784 27360 21784 10 10 3
1 pool-1 256 RAID5 NORMAL 2656 3168 2656 10 DIF 9 3
2 pool-2 64 RAID6 NORMAL 1248 2112 1248 10 DIF 6 3
3 pool-3 32 RAID5 NORMAL 1248 1760 1248 10 DIF 5 3
4 pool-4 256 RAID6 NORMAL 3040 4000 3040 10 DIF 10 3
5 pool-5 32 RAID6 NORMAL 1440 2400 1440 10 DIF 6 3
6 pool-6 128 RAID6 NORMAL 3040 4000 3040 10 DIF 10 3
Select capacity:
MAX
The Virtual Disk capacity will be rounded up to an 8GB boundary.
(NOTE: If MAX is typed then the maximum free capacity available in each pool will be used)
CREATE VIRTUAL DISK

4.7.3 Configure Virtual Disk Attributes

To change a virtual disk's attributes, select **Virtual Disk > Set Virtual Disk Attributes**.

At the Set Attributes for Pool screen (Figure 145), select the VD whose attributes you want to change. Make the desired changes. Click **UPDATE** to save the changes.

Figure 145. Set Virtual Disk Attributes Screen

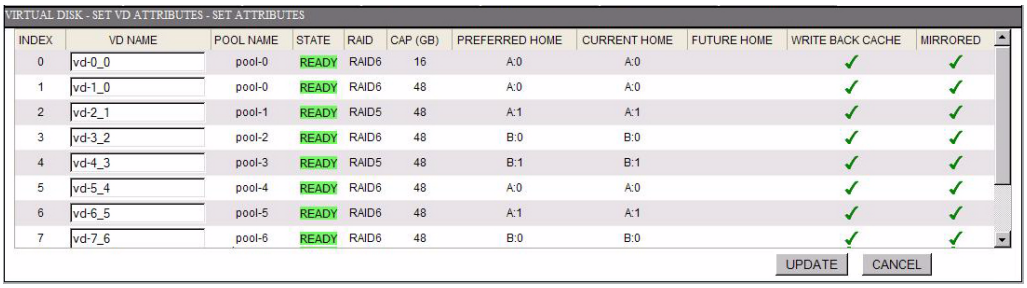
VIRTUAL DISK - SET VD ATTRIBUTES - SET ATTRIBUTES
INDEX
VD NAME
POOL NAME
STATE
RAID
CAP (GB)
PREFERRED HOME
CURRENT HOME
FUTURE HOME
WRITE BACK CACHE
MIRRORED
0 vd-0_0 pool-0 READY RAID6 16 A:0 A:0
1 vd-1_0 pool-0 READY RAID6 48 A:0 A:0
2 vd-2_1 pool-1 READY RAID5 48 A:1 A:1
3 vd-3_2 pool-2 READY RAID6 48 B:0 B:0
4 vd-4_3 pool-3 READY RAID5 48 B:1 B:1
5 vd-5_4 pool-4 READY RAID6 48 A:0 A:0
6 vd-6_5 pool-5 READY RAID6 48 A:1 A:1
7 vd-7_6 pool-6 READY RAID6 48 B:0 B:0
UPDATE
CANCEL

4.7.4 Delete a Virtual Disk

Warning

When you delete a Virtual Disk, you lose access to all data that was stored on that Virtual Disk.

If you need to delete a virtual disk, select **Virtual Disk > Delete Virtual Disk**.

At the Delete Virtual Disk screen (Figure 146), select the virtual disk that you want to delete.Click **DELETE** to delete that VD or click **CANCEL** to cancel operation. You will be prompted to confirm deletion. Click **OK** to confirm.

NOTE : If a presentation is configured for a virtual disk, deleting the virtual disk will return an error. You must first delete all the presentations configured on that virtual disk using the **Delete Presentation** function under the **Presentation** menu.

Figure 146. Delete Virtual Disk Screen

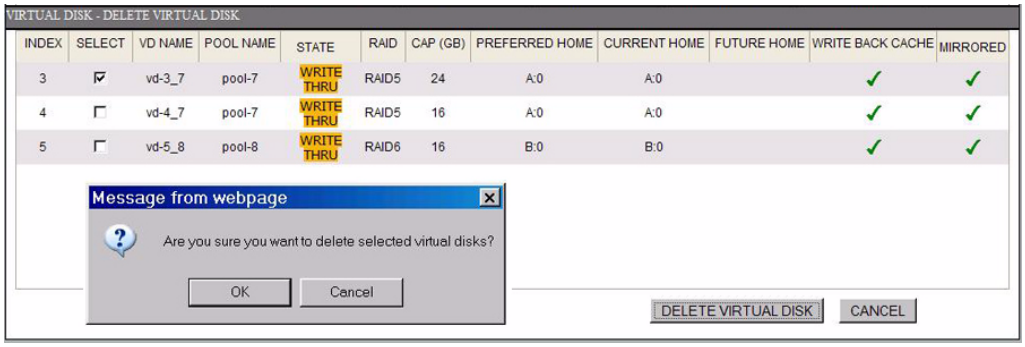

4.8 Spare Pools

The SFA10000 supports the concept of spare pool. A spare pool contains physical disks that can be used as spare disks in one or more storage pools. In the event of a disk failure, a disk that is assigned to a spare pool is automatically swapped into the storage pool to replace the failed disk. A rebuild then occurs immediately, reducing the possibility of data loss due to additional disk failures.

Spare pools can be shared by storage pools or dedicated to a particular storage pool in which case the spare disks will only be used by the designated storage pool.

NOTE : Each storage pool must have a spare pool assigned to it.

The spare pool commands are located under the **Pools** pull down menu (Figure 147).

Figure 147. Pools Menu

4.8.1 Show Spare Pools

To view the list configured spare pools, select **Pools > Show Spare Pools**.

Click on the individual spare pool name to display its detailed information (Figure 148). Click on the **Total Member** column to display the list of disks in that spare pool.

Figure 148. Show Spare Pool Screen

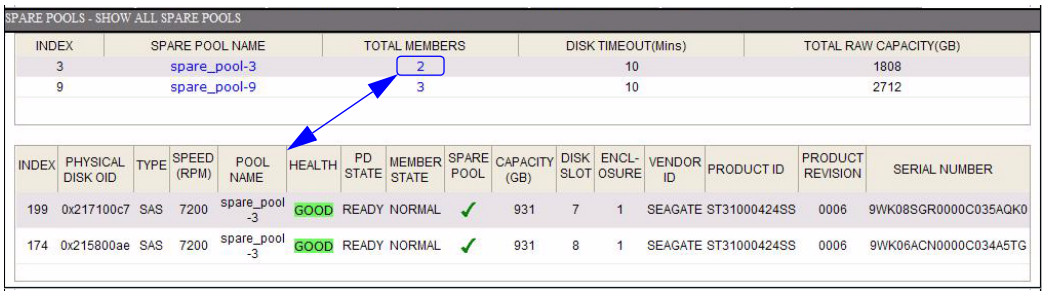

4.8.2 Locate a Spare Pool

To locate a spare pool, select **Pools > Locate Spare Pools**.

At the Locate Spare Pool screen (Figure 149), select the spare pool and the screen changes accordingly showing the disk locations for that pool. Click **LOCATE** to flash the status LED of those disks.

Figure 149. Locate Spare Pool

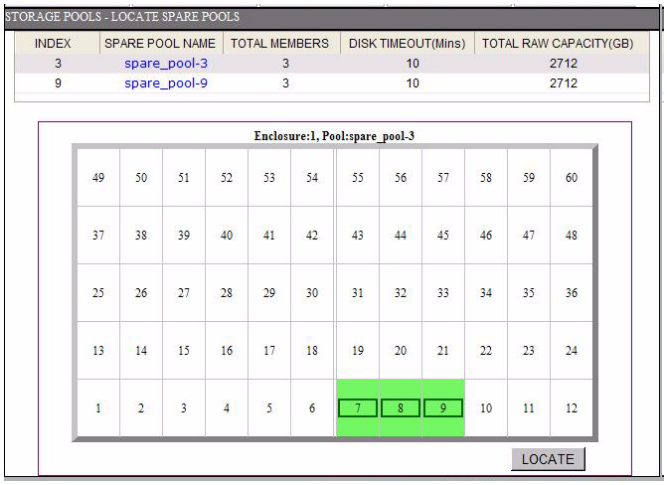

4.8.3 Create a Spare Pool

To create a spare pool, select **Pools > Create Spare Pool**.

At the Create Spare Pool screen (Figure 150), select the drive size, spindle speed, and drive type for the pool.

Figure 150. Create Spare Pool

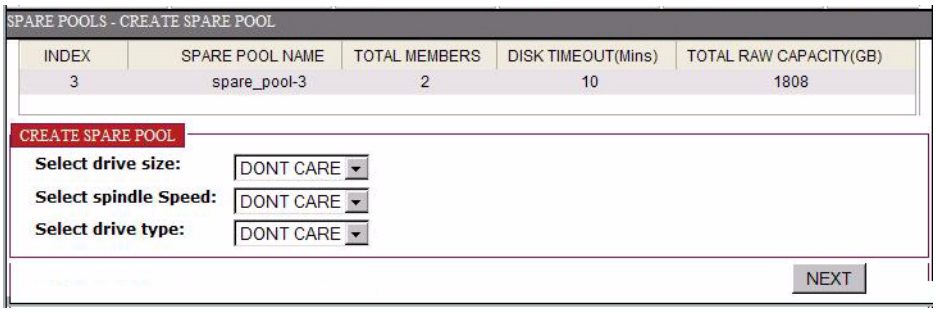

Click **NEXT** to bring up the Create Pool Advanced screen (Figure 151). Check to select the disks to add to the pool to be created. Click **CREATE SPARE POOL** to create the pool.

NOTE : When assigning disk drives to a spare pool that is to be assigned to a given storage pool, be sure to select a disk that is as large or larger than any disks in the assigned storage pool.

Figure 151. Select Disks for Spare Pool

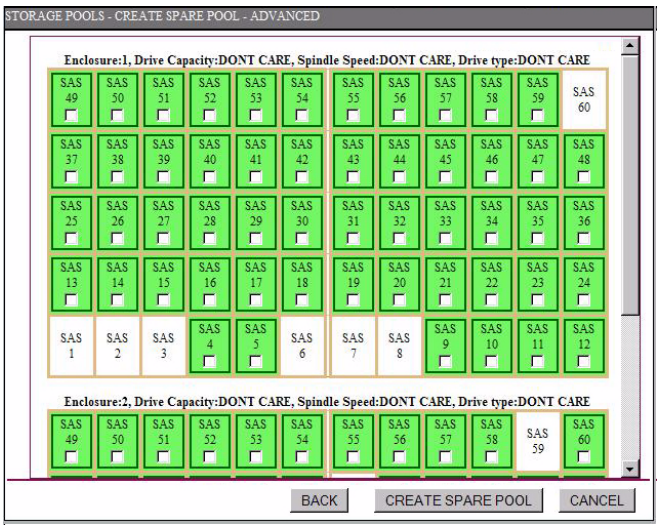
STORAGE POOLS - CREATE SPARE POOL - ADVANCED
Enclosure:1, Drive Capacity:DONT CARE, Spindle Speed:DONT CARE, Drive type:DONT CARE
Enclosure:2, Drive Capacity:DONT CARE, Spindle Speed:DONT CARE, Drive type:DONT CARE
BACK
CREATE SPARE POOL
CANCEL

4.8.4 Assign Spare Pool to Storage Pool

To assign a spare pool to a storage pool, select **Pools > Set Spare Pool Drive Assign Policy**.

At the Spare Pool Drive Assign Policy screen (Figure 152), make the desired changes and click **UPDATE** to save the changes.

NOTE : Each storage pool should have a spare pool assigned to it. If the Select Spare Pool entry shows "NONE", no spare pool will serve this storage pool.

Figure 152. Assign Spare Pool to Storage Pool Screen

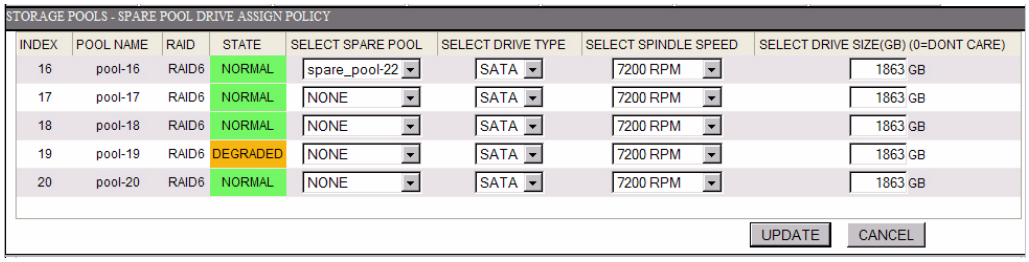
STORAGE POOLS - SPARE POOL DRIVE ASSIGN POLICY
INDEX
POOL NAME
RAID
STATE
SELECT SPARE POOL
SELECT DRIVE TYPE
SELECT SPINDLE SPEED
SELECT DRIVE SIZE(GB) (0=DONT CARE)
16 pool-16 RAID6 NORMAL spare_pool-22 SATA 7200 RPM 1863 GB
17 pool-17 RAID6 NORMAL NONE SATA 7200 RPM 1863 GB
18 pool-18 RAID6 NORMAL NONE SATA 7200 RPM 1863 GB
19 pool-19 RAID6 DEGRADED NONE SATA 7200 RPM 1863 GB
20 pool-20 RAID6 NORMAL NONE SATA 7200 RPM 1863 GB
UPDATE
CANCEL

4.8.5 Set Spare Pool Attributes

To change the spare pool's attributes, select **Pools > Set Spare Pool Attributes**.

At the Set Attributes for Pool screen (Figure 153), click to select the spare pool whose attributes you want to change. Make the desired changes. Click **UPDATE** to save the changes.

Figure 153. Set Spare Pool Attributes Screen

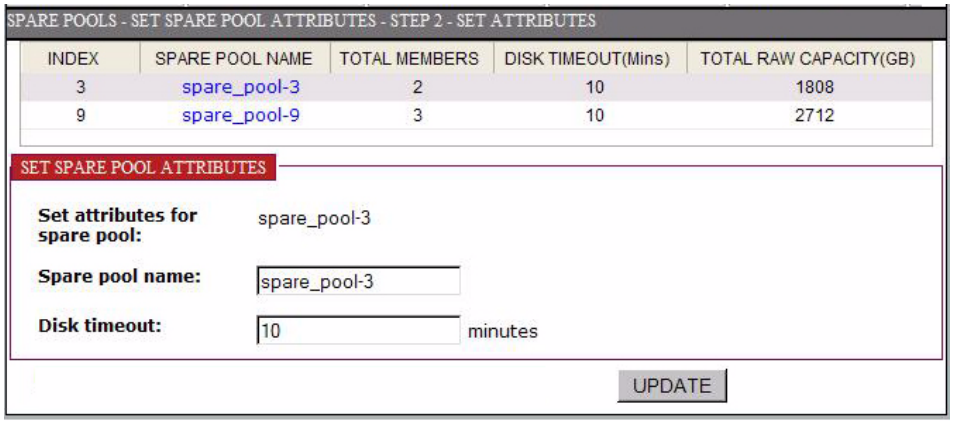

4.8.6 Delete a Spare Pool

In you need to delete a spare pool, select **Pools > Delete Spare Pool**.

At the Delete Spare Pool screen (Figure 154), select the spare pool you want to delete. Click **DELETE SPARE POOLS**. You will be prompted to confirm deletion. Click **OK** to proceed or CANCEL to cancel operation.

Figure 154. Delete Set Spare Pool Screen

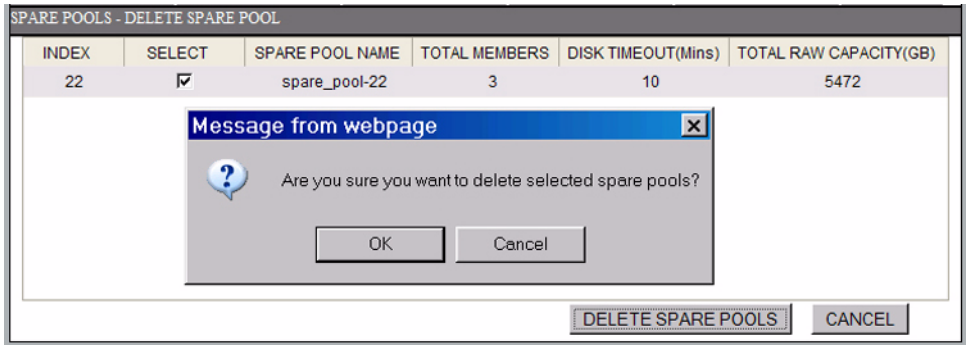

4.8.7 Delete a Disk from Spare Pool

If you need to remove a disk from a spare pool, refer to Section 4.4.3, "Remove Disk from Spare Pool" on page 116.

4.9 Presentations

NOTE : For SFA10000E, this feature is not applicable.

Virtual disks are only presented to the hosts that have been given authorized access. A Presentation on a SFA10000 has the following components:

- Discovered Initiators
  A Discovered Initiator is defined as a Fibre Channel or InfiniBand initiator that has logged into the SFA10000. This is normally a single port on a Fibre Channel or InfiniBand HBA in an external computer.
- Host
  A Host is an abstract client of a virtual disk that has all the access to all data within the virtual disk. This most commonly represents an external computer with one or more Fibre Channel or InfiniBand initiators. Host attributes are OS (operating system) of which you can select GENERIC, WINDOWS, LINUX or MAC_OS. The OS attribute dictates the way the controller responds to certain Fibre Channel commands where appropriate. For example, Mac OS X expects a different response from an inquiry of LUN0 (the controller LUN) than a Windows host.
- Channel
  A Channel is one of the Fibre Channel or InfiniBand ports on the controller. Channel attributes are MODE of which you can select either MAC_OS or STANDARD. When MAC_OS is selected, the port's node name will be set differently in order to be visible to a Macintosh system.
- Stack
  A Stack is the input side of the subsystem. In the case of the SFA10000, the stack is defined as the Fibre Channel interface. In future versions of the product, different interface stacks may exist.
- Presentation
  A Presentation is the relationship between a Host and a virtual disk. A Presentation implies that the related Host has some sort of access to the virtual disk. Attributes of a Presentation are:
  - PORT: from which the host will see the virtual disk
  - READ_ONLY: controls read only access
  - PRESENT_HOME_ONLY: presents the specified virtual disk from its designated home controller only
  - LUN: user-specified LUN number that the virtual disk will show to the host.

4.9.1 Presentation Menu

Figure 155 shows the **Presentation** pull down menu.

Figure 155. Presentation Menu

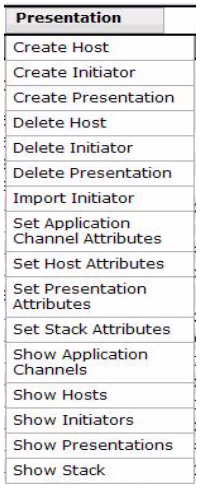

4.9.2 Show Presentations

To display the list of configured presentations (Figure 156), select **Presentation > Show Presentations**.

Figure 156. Show Presentations Screen

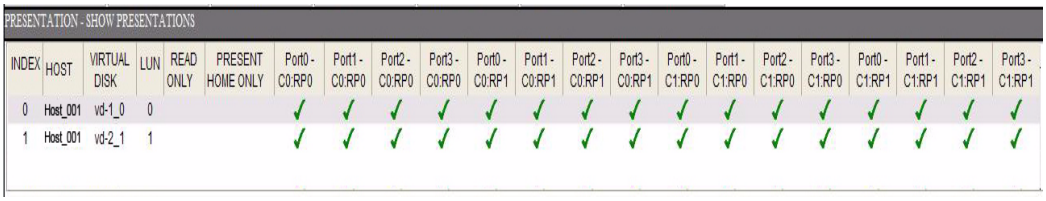

4.9.3 Set Up a Presentation

4.9.3.1 Create a Host

1. Select **Presentation > Create Host**.
2. At the Create Host screen (Figure 157), enter a host name and specify the Stack and OS attributes. Click **CREATE HOST**.

Figure 157. Create Host Screen

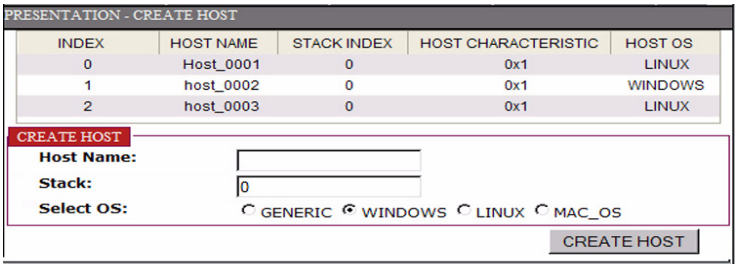

4.9.3.2 Map Host to Discovered Initiators

1. Select **Presentation > Import Initiators**.
2. At the Import Initiator screen (Figure 158), select an initiator and the corresponding host(s).
3. Click **IMPORT INITIATOR**.

Figure 158. Import Initiator Screen

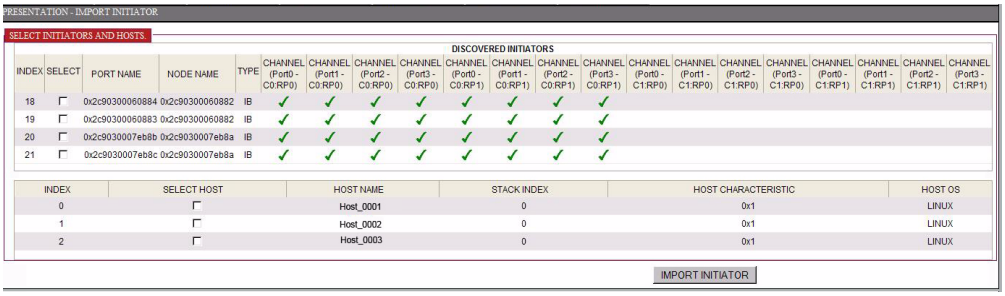

4.9.3.3 Present Virtual Disks to Hosts

1. Select **Presentation > Create Presentation**.
2. At the Create Presentation screen (Figure 159), select a host and the virtual disk(s).

Figure 159. Create Presentation Screen

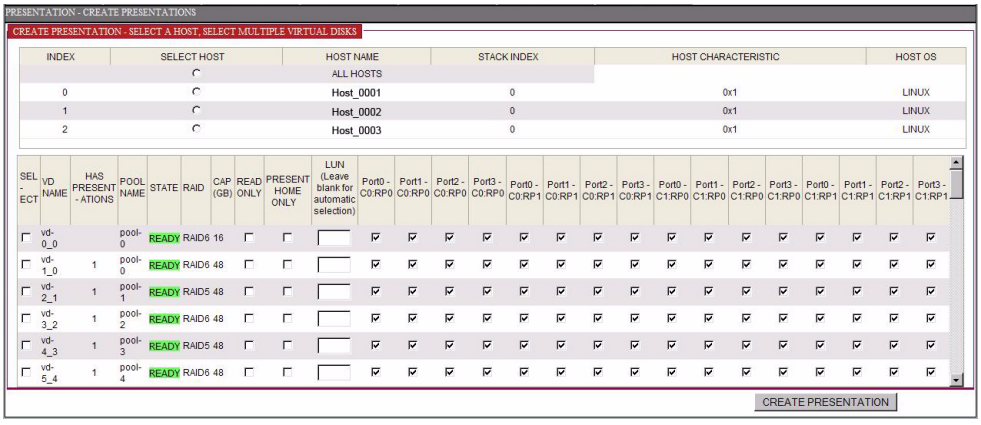

OPTIONAL: To present a virtual disk to all host ports on both controllers (promiscuous mode), select “ALL HOSTS” instead of individual hosts.

NOTE : The promiscuous mode may not be used with persistent reservations. The “ALL HOSTS” setting will cause SCSI2/3 Reservation commands to fail.

3. Then select the mask options and enter the Logical Unit Number (LUN).
4. Click **CREATE PRESENTATION**. The newly created presentation is displayed.

4.10 Enclosure Menu

Figure 160 shows the **Enclosures** pull down menu.

Figure 160. Enclosure Menu

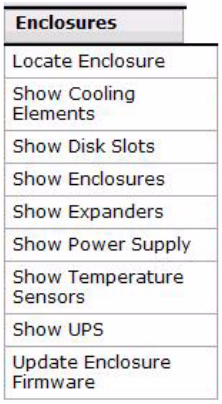

4.10.1 Show Enclosures

The **Show Enclosures** function displays all the controllers and disk enclosures that are installed in the SFA10000 system (Figure 161). Clicking on the highlighted item will bring up the detailed information of that component.

Figure 161. Show Enclosures Screen

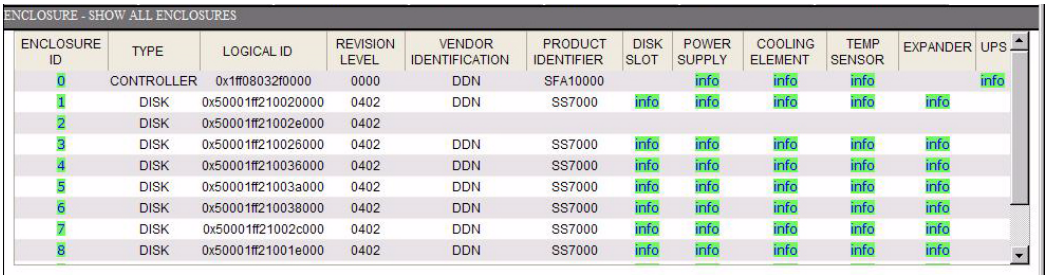

ENCLOSURE ID	TYPE	LOGICAL ID	REVISION LEVEL	VENDOR IDENTIFICATION	PRODUCT IDENTIFIER	DISK SLOT	POWER SUPPLY	COOLING ELEMENT	TEMP SENSOR	EXPANDER	UPS
0	CONTROLLER	0x1ff08032f0000	0000	DDN	SFA10000		info	info	info		info
1	DISK	0x50001ff210020000	0402	DDN	SS7000	info	info	info	info	info	
2	DISK	0x50001ff21002e000	0402								
3	DISK	0x50001ff210026000	0402	DDN	SS7000	info	info	info	info	info	
4	DISK	0x50001ff210036000	0402	DDN	SS7000	info	info	info	info	info	
5	DISK	0x50001ff21003a000	0402	DDN	SS7000	info	info	info	info	info	
6	DISK	0x50001ff210038000	0402	DDN	SS7000	info	info	info	info	info	
7	DISK	0x50001ff21002c000	0402	DDN	SS7000	info	info	info	info	info	
8	DISK	0x50001ff21001e000	0402	DDN	SS7000	info	info	info	info	info	

4.10.2 Update Enclosure Firmware

You can update the disk enclosure firmware using the **Update Enclosure Firmware** function. After you have obtained the new firmware files, select **Enclosures > Update Enclosure Firmware**.

Figure 162. Update Enclosure Firmware Screen

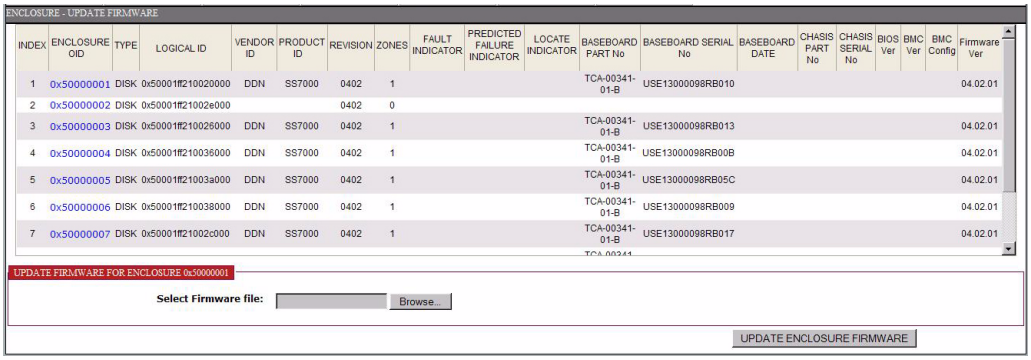
ENCLOSURE - UPDATE FIRMWARE
INDEX
ENCLOSURE OID
TYPE
LOGICAL ID
VENDOR ID
PRODUCT ID
REVISION
ZONES
FAULT INDICATOR
PREDICTED FAILURE INDICATOR
LOCATE INDICATOR
BASEBOARD PART No
BASEBOARD SERIAL No
BASEBOARD DATE
CHASIS PART No
CHASIS SERIAL No
BIOS Ver
BMC Ver
BMC Config
Firmware Ver
1 0x50000001 DISK 0x50001ff210020000 DDN SS7000 0402 1 TCA-00341-01-B USE13000098RB010 04.02.01
2 0x50000002 DISK 0x50001ff21002e000 0402 0
3 0x50000003 DISK 0x50001ff210026000 DDN SS7000 0402 1 TCA-00341-01-B USE13000098RB013 04.02.01
4 0x50000004 DISK 0x50001ff210036000 DDN SS7000 0402 1 TCA-00341-01-B USE13000098RB00B 04.02.01
5 0x50000005 DISK 0x50001ff21003a000 DDN SS7000 0402 1 TCA-00341-01-B USE13000098RB05C 04.02.01
6 0x50000006 DISK 0x50001ff210038000 DDN SS7000 0402 1 TCA-00341-01-B USE13000098RB009 04.02.01
7 0x50000007 DISK 0x50001ff21002c000 DDN SS7000 0402 1 TCA-00341-01-B USE13000098RB017 04.02.01
UPDATE FIRMWARE FOR ENCLOSURE 0x50000001
Select Firmware file:
Browse...
UPDATE ENCLOSURE FIRMWARE

At the Update Firmware screen (Figure 162), click to select the enclosure to update. Locate the firmware file and click **UPDATE ENCLOSURE FIRMWARE**. Then follow the onscreen instructions to complete the update. Also refer to Section 3.11.3, "Disk Enclosure Firmware Update Procedure" on page 92 for detailed instructions.

Chapter 5

Support

This section describes common problems, with possible solutions, which can occur with the SFA10000 system.

5.1 "RAID[0]$" Prompt is not shown correctly

The default CLUI prompt is "RAID[0]$". If "CLUI$" is shown instead of "RAID[0]$", one of the following problems may have occurred:

- The firmware on this controller is in a "MIR" state and the MIR state needs to be resolved. Refer to Section 5.3, "Manual Intervention Required (MIR) States" on page 137 for troubleshooting information.
- The CLUI has lost communication with the firmware or the firmware is not running. A restart of the system is needed.
- There are multiple CLUI connections to the firmware already and there are not enough resources to allow another connection. Only one CLUI connection is currently supported.

5.2 Event Log

Some types of issues can be seen from the event log:

- Background jobs affecting performance

```
014547 2009-08-31 12:30:08:1953906  G=62 S=0 T=1 RP=0 VP=1 LOG_ST_MEMBER_CHANGE
STATE POOL MEMBER CHANGE INDEX:00000000 ID:0x5000cca20def4516 PDIDI:0085 MEMBER_INDEX:0000
STATE:REBUILD
```

- Backend issues affecting performance

```
000425 2009-04-28 14:39:12:6963749  G=10 S=0 T=1 RP=0 VP=5
LOG_AMPD_MPI_SAS_DEVICE_DISAPPEARED   AMPD IOC0 WWN:5000cca216ed8430
ES:50001ff101ed0000:60. RC:04 Hndl:0065 PrntId:50001ff101ed017f PHY:08
BT:00:14 DI:00000081 DSK:0000

000430 2009-04-28 14:39:13:8133330  G=10 S=0 T=1 RP=0 VP=5
LOG_AMPD_MPI_SAS_DEVICE_APPEARED   AMPD IOC0 WWN:5000cca216ed8430
ES:50001ff101ed0000:60. RC:03 Hndl:0065 PrntId:50001ff101ed017f PHY:08
BT:00:14 DI:00000081 DSK:0000
```

- ICL problem - link bounce

```
000468 2009-08-25 22:04:20:7729086  G=60 S=0 T=1 RP=0 VP=5 LOG_AMPD_ICL_LINK_STATUS_CHANGE
AMPD Ctx:2f0006c24528 Drv:2f0006c2ed80 Old:DWN New:UP
000469 2009-08-25 22:04:20:7729497  G=60 S=0 T=1 RP=0 VP=1 LOG_DUCK_CONTROLLER_CONNECTED
DUCK Controller Connected.
```

- ICL problem - controller failure

```
008418 2009-07-29 08:16:50:9043234  G=60 S=0 T=1 RP=0 VP=1 LOG_DUCK_CONTROLLER_DISCONNECTED
DUCK Controller Disconnected.
008419 2009-07-29 08:16:50:9043363  G=60 S=0 T=1 RP=0 VP=1 LOG_ST_OTHER_DIED
STATE OTHER CONTROLLER DIED

RAID[0]$ show controller *
OID: 0x38000000 Index: 0000 Name: A LOCAL PRIMARY
Total Controllers: 1
```

Refer to the **DDN SFA Service Manual** for a complete list of event log entries.

5.3 Manual Intervention Required (MIR) States

Manual Intervention Required (MIR) represents a condition with the controller that requires the user to provide a solution before proceeding with normal controller operations. This is to guard against the controller firmware from executing operations that may not necessarily be the desired operation of the user. These conditions will most likely be seen in a new system installation environment. For example, when a system is booted and the backend physical disks have never been installed behind the controllers, the controller firmware has never had a chance to write out configuration metadata. The controller recognizes that there is no valid metadata and requires the user to acknowledge proceeding or not.

Use the **SHOW CONTROLLER LOCAL ALL_ATTRIBUTES** command to display the MIR condition of the controller you are logged into (Figure 163).

Figure 163. Show Controller Information Screen

```
RAID[0]$ show controller 0 all
Index:                    0
OID:                      0x38000000
Firmware Version:
  Release:                1.4
  Source Version:         4476
  Fully Checked In:       Yes
  Private Build:          No
  Build Type:             Production
  Build Date and Time:    2010-05-28-22:10UTC
  Builder Username:       root
  Builder Hostname:       co-bs2
  Build for CPU Type:     AMD-64-bit
Hardware Version:         0000
State:                    RUNNING
 .
 .
 .
Name:                     A
Controller:               LOCAL        (SECONDARY)
Controller ID:            0x0001ff0800a30000
Enclosure OID:            0x50000006 (Index 6)
Universal LAN Address:    0x00000001ff0800a3
MIR Reason:               None
```

Listed below are some commonly seen MIR conditions and their required actions. Refer to the *DDN SFA OS Service Manual* for a complete list of MIR conditions.

MIR_JIS_DISCOVERY_IN_PROG

Initialized Storage discovery is in progress. Please allow time for configuration discovery to complete. If this condition persists, please reboot the system. If this condition continues to persist after the reboot, please contact customer support.

MIR_OTHER_JIS_DISCOVERY_IN_PROG

Initialized Storage discovery is in progress on other controller. Please allow time for configuration discovery to complete. If this condition persists, please reboot the system. If this condition continues to persist after the reboot, please contact customer support.

MIR_NO_BACKEND_DRIVES

This controller cannot find any disk modules on the backend. Install disk modules on the backend or fix the condition that prevents this controller from finding backend disks. Please refer to the installation guidelines for proper setup.

MIR_NO_CONFIG

A configuration could not be created on the backend disks. Fix the condition that prevents this controller from creating a configuration on the backend disks. Please refer to the installation guidelines for proper setup.

MIR_NO_QUORUM

No quorum disks could be found within the disk modules on the backend. Fix the condition that prevents the quorum disks from being seen. This may be caused by the disks never being in the storage array. This will normally happen on new system installations. Use the "**CLEAR SUBSYSTEM MIR_STATE**" command to create an empty configuration. ***NOTE**: This creates a new configuration and the old configuration will be deleted if existed.

MIR_NOT_LAST_CONTROLLER

This controller found a valid configuration, but was not present when another controller owned the configuration. Since the controller was not present, it does not know if another controller may have cached data for this configuration. This may happen when controllers are swapped out or if the controllers went down and restarted individually. In many cases, this MIR state will resolve itself by communicating with the other controller and obtaining permission to use the found configuration. If the MIR state does not resolve itself in 3 minutes, contact DDN Technical Support.

MIR_MULTIPLE_JIS

Multiple configurations were found on the quorum disks. This may happen if the disks from one system were installed in another system when the systems were powered down and then rebooted. If the disks were installed in the system while running, then this should not be an issue. A list of found configurations will be listed. Use the "**CLEAR SUBSYSTEM MIR_STATE ID=<id>**" command to use the specified ID's configuration.

MIR_DUAL_NO_AGREE

The two connected controllers do not agree on the ID of the configuration. This may occur if one controller saw a subset of the disks and the other controller saw a different subset of disks. Please refer to the installation guidelines for proper setup.

MIR_CONFIG_MISMATCH

The configuration version of this firmware does not match the configuration version of that on media. To proceed, either reload the previous version of firmware and do a backup then upgrade, or delete your configuration to continue. Use "**CLEAR SUBSYSTEM CONFIGURATION**" to create an empty configuration. ***NOTE**: This creates a new configuration and the old configuration will be deleted if existed.

MIR_NO_LOAD_CONFIG

A configuration could not be loaded from the backend disks. Fix the condition that prevents this controller from loading a configuration from the backend disks, or use "**CLEAR SUBSYSTEM MIR_STATE**" to create an empty configuration. Please refer to the installation guidelines for proper setup. ***NOTE**: This creates a new configuration and the old configuration will be deleted if existed.

5.4 Recovery from Disk Failure

5.4.1 Automatic Rebuild

A single disk failure in any storage pool does not result in data loss. The virtual disk(s) on that storage pool will continue to operate in degraded mode. If a spare disk is available and automatic rebuild is enabled, the SFA10000 will automatically rebuild the data to a spare disk.

NOTE : System performance will be impacted while recovery is taking place.

When a disk failure occurs, the failure is written to the event log. You may monitor the rebuild progress or adjust the rebuild rate to match the user load requirements.

To display the event log (Figure 164), use the **SHOW CONTROLLER LOCAL LOG** command.

Figure 164. Display Event Log

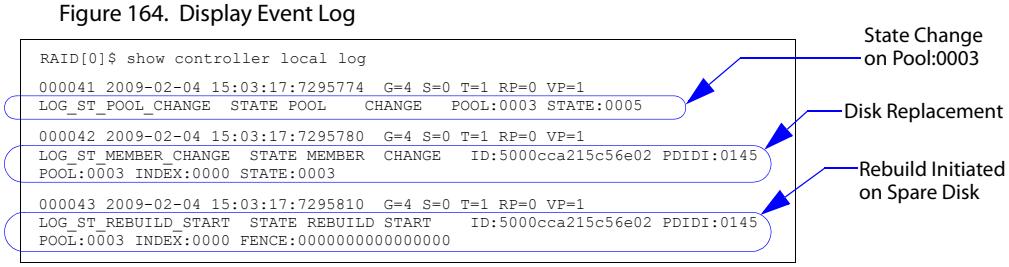
RAID[0]$ show controller local log
000041 2009-02-04 15:03:17:7295774 G=4 S=0 T=1 RP=0 VP=1
LOG_ST_POOL_CHANGE STATE POOL CHANGE POOL:0003 STATE:0005
000042 2009-02-04 15:03:17:7295780 G=4 S=0 T=1 RP=0 VP=1
LOG_ST_MEMBER_CHANGE STATE MEMBER CHANGE ID:5000cca215c56e02 PDIDI:0145
POOL:0003 INDEX:0000 STATE:0003
000043 2009-02-04 15:03:17:7295810 G=4 S=0 T=1 RP=0 VP=1
LOG_ST_REBUILD_START STATE REBUILD START ID:5000cca215c56e02 PDIDI:0145
POOL:0003 INDEX:0000 FENCE:0000000000000000
State Change on Pool:0003
Disk Replacement
Rebuild Initiated on Spare Disk

To look at the failed disk (Figure 165), enter command **SHOW UNASSIGNED_POOL FAILED ALL**

Figure 165. Display Failed Disk

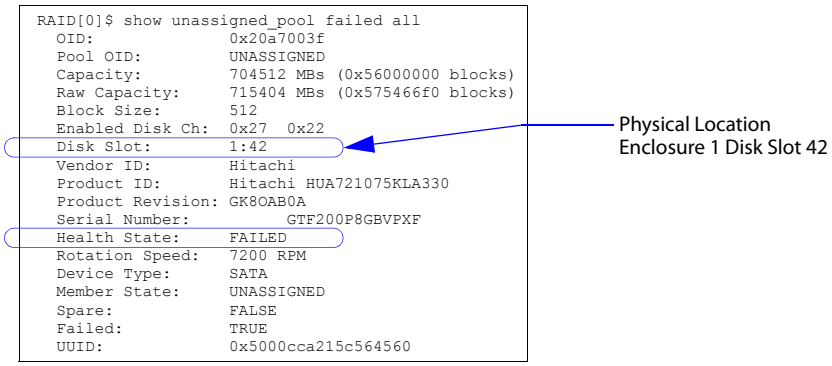

To monitor the rebuild progress (Figure 166), enter command **SHOW JOB * ALL_ATTRIBUTES**

Figure 166. Display Rebuild Progress

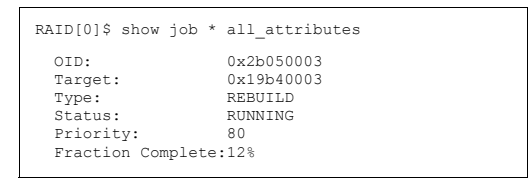

To show the storage pool information (Figure 167), enter command **SHOW POOL <id> ALL_ATTRIBUTES**. Once the rebuild is complete, the status of storage pool will return to "NORMAL".

Figure 167. Display Storage Pool Information

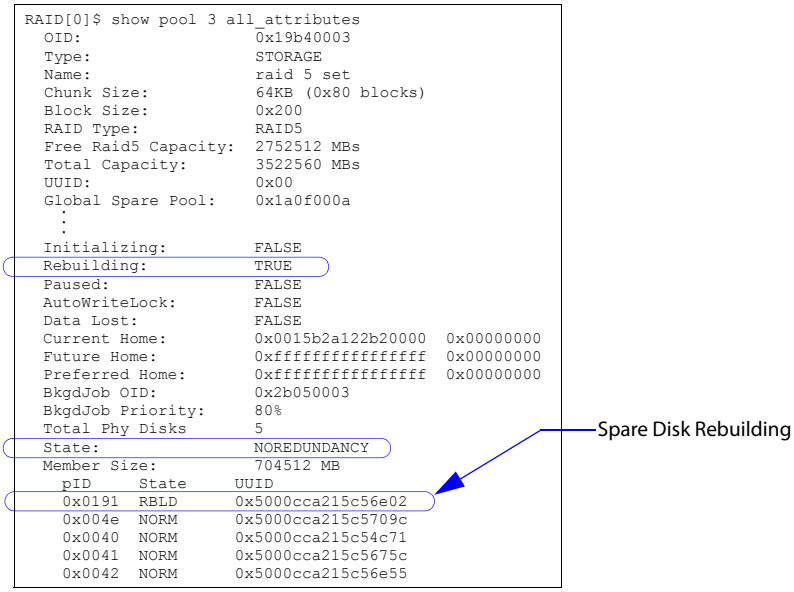

5.4.2 Spare Disk Replacement

Upon completion of a rebuild, the spare disk becomes a member of the storage pool, replacing the failed disk. After you have replaced the failed disk with a new disk, the new disk is added to the "Unassigned Pool". It is recommended that you assign this new disk to the spare pool to "replace" the spare disk that has been used.

5.4.3 When a Spare is not available

When a disk is failed by the system and there is no spare disk available, you need to replace the failed disk immediately. After you have replaced the failed disk with a new disk, you can initiate a rebuild as described below:

1. Enter command **SHOW UNASSIGNED_POOL * ALL** to identify the new disk's index name.

2. Enter command **ASSIGN PHYSICAL_DISK <new-disk-id> TO_POOL <pool-id> SET_SPARE** where <new-disk-id> is the index name of the replacement disk and <pool-id> is the OID of the storage pool that had the failed disk.

5.4.4 Manual Rebuild

You may manually replace a failed disk using the **REPLACE** command:

REPLACE PHYSICAL_DISK <id> NEW_DISK <new-disk-id>

A Replace operation is used to replace a failed disk with a healthy spare disk. The operation can take several hours to complete depending on the size of the disk and speed of the replace operation. The rate of rebuild can be adjusted (see Section 3.11, "Firmware Update Management" on page 88 for more information).

5.4.5 Power Cycling Disk Drives

SATA disks can get wedged and stop working. When this occurs, power cycling may bring the disk back on-line.

To power cycle a disk, enter the CLUI command:

```
SET SLOT=<enclosure-id> <slot-id> POWER
```

NOTE : Disk power cycling requires an interposer that supports power cycling.

Whether a disk can be power cycled depends on the following factors:

- Enclosure
  - Enclosure type
  - Enclosure firmware version
- Interposer (Multiplexor)
  - Interposer type
  - Interposer hardware version

Table 6 shows the current enclosures and interposers along with whether they can be power cycled or not.

Table 6. Disk Slot Power Cycle Support

Enclosure	Interposer	Power Cycle Supported
SS6000	SAS-to-SAS	Yes
	SATA Multiplexor	Yes
	SAS-to-SATA Bridge	No
	None (SAS only)	No
SS7000	SAS-to-SAS	No
	SAS-to-SATA Bridge	No, but
	None (SAS only)	No
SS2460	None	No

5.5 UPS Maintenance

This section explains how to care for the SFA10000's UPS and battery pack, how to replace the battery pack, and how to test the battery pack.

5.5.1 UPS and Battery Pack Care

The UPS contains a battery pack consisting of three batteries, a frame, and wiring. The batteries have a useful life of between two and three years. To insure that the UPS will operate when a power failure occurs, the UPS battery pack should be replaced every 24 months even if the battery pack still tests okay.

If you store the UPS or a battery pack for a long period, recharge the battery pack every six months by plugging the UPS into a power outlet. To recharge a battery pack, put it in a UPS and then plug the UPS into a power outlet. The batteries charge to 90% capacity in approximately three hours. However, it is recommended that battery packs be charged for 6 to 24 hours after long-term storage.

5.5.2 Checking the UPS Status

The **SHOW UPS ALL_ATTRIBUTES** command shows the battery manufacture date and the battery life remaining in days (Figure 168).

Figure 168. Show UPS Attributes

```
******************
*     UPS(s)     *
******************
Sub Index:                1
Sub OID:                  0x80000001
Enclosure Index:          0
Enclosure OID:            0x50000000
Position:                 1
SES Status:               OK
Charge Level:             100 %
Hold-Up Time Remaining:   17 minutes
Enabled:                  TRUE
AC Failed:                FALSE
Present:                  TRUE
Interface Failed:         FALSE
Battery Health            OK
Fault Indicator:          OFF
Predicted Failure Ind:    OFF
Locate Indicator:         OFF
Battery Mfg. Date:        Fri Nov12 12:55:43 2010
Battery Life Remaining:   365 days
Part number:              PS5115 RM 1500VA
Serial number:            GD371A0853
Firmware version:         Cont: 01.20 Inve:01.02
```

When the controller starts up, ES checks the battery manufacture date against the current time and the expected battery life. A battery end of life warning will be issued as the battery reaches the specified dates:

- If the battery life remaining is between 0 and 30 days, it will log a warning.

```
000483 2010-11-12 12:53:47:8178400  G=62 S=1 T=1 RP=0 VP=01
LOG_ES_BATTERY_END_OF_LIFE_NEAR     ES Controller 0xc2999211d0001 UPS
battery at position 1 is nearing its end-of-life and should be replaced
within 30 days.
```

- If the battery life remaining is 0 or less it will log an error.

```
000486 2010-11-12 12:54:52:3214473  G=61 S=2 T=1 RP=0 VP=01
LOG_ES_BATTERY_END_OF_LIFE_PASSED    ES Controller 0xc2999211d0001 UPS
battery at position 1 has reached its end-of-life (0 days ago) and should
be replaced immediately.
```

In both cases, the battery life remaining is reported "…should be replaced within N days." Note that if the battery replacement is overdue then the error will read as something like "…(10 days ago)…".

These events will also be logged while the controller is running as they cross the thresholds at 30 and 0 days of battery life remaining. Note that there will be two other events that are logged with these which are logged any time the battery changes state. These events should occur on the order of once a year so there should be limited traffic to the event log. Note that `LOG_ES_ENCL_UPS_WARN_BATT` may occur for other reasons than the battery reached its end of life.

```
000486 2010-11-12 12:54:52:3214473  G=61 S=2 T=1 RP=0 VP=01
LOG_ES_BATTERY_END_OF_LIFE_PASSED    ES Controller 0xc2999211d0001 UPS
battery at position 1 has reached its end-of-life (0 days ago) and should
be replaced immediately.

000487 2010-11-12 12:54:52:3214503  G=60 S=0 T=1 RP=0 VP=01
LOG_ES_ENCL_UPS_INFO     ES UPS for controller 0xc2999211d0001 changed
state SES status = 2, disabled = 0, charge = -1%, hold-up time = -2 min.

000488 2010-11-12 12:54:52:3214512  G=61 S=1 T=1 RP=0 VP=01
LOG_ES_ENCL_UPS_WARN_BATT     ES UPS for controller 0xc2999211d0001
indicates that its battery should be replaced. SES(2 0 0 8).
```

5.5.3 Replacing the Battery Pack

NOTE : Do not disconnect the battery pack while the UPS is in Battery mode.

CAUTION ! If there is a power outage while replacing the battery pack, the controller without its battery pack will power down. However, the redundant controller insures that no operations are lost or corrupted. For this reason, when you change the battery pack on one controller, let it charge completely before changing the battery pack on the other controller.

With the hot-swappable battery feature, the UPS battery packs can be replaced without turning the UPS off or disconnecting the load. Note that the battery packs are the FRU, not the batteries.

If you prefer to shut down the UPS to change the battery pack, use the CLUI to issue a **`SHUTDOWN CONTROLLER LOCAL`** to the corresponding controller. This will cause an orderly controller shutdown which concludes by automatically turning off the UPS power.

WARNING:
Batteries can present a risk of electrical shock or burn from high short circuit current. The following precautions should be observed: 1) Remove watches, rings, or other metal objects; 2) Use tools with insulated handles; 3) Do not lay tools or metal parts on top of batteries.

ELECTRIC ENERGY HAZARD. Do not attempt to alter any battery pack wiring or connectors. Attempting to alter wiring can cause injury.

Replace the battery pack with the same type of battery pack as originally installed in the UPS.

Perform the following steps to replace the battery pack:

NOTE : Replace the battery pack every 24 months even if it is still functioning.

1. Remove the UPS bezel by loosening the right side and then tapping the right side to disengage the left side (Figure 169).

Figure 169. Removing the UPS Bezel

2. Disconnect the battery pack connector (Figure 170).

Figure 170. Disconnecting the Battery Pack Connector

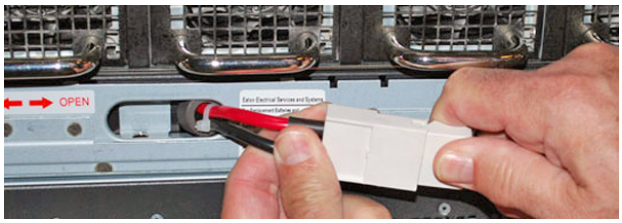

3. Slide the battery pack cover to the right and pull it off (Figure 171). Set aside the battery pack cover.

Figure 171. Removing the Battery Pack Cover

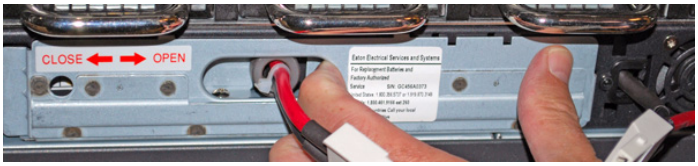

4. Pull the battery pack out and place it on a flat, stable surface (Figure 172).

CAUTION ! Note that the battery pack weighs 20 pounds. Be sure to support the battery pack as you pull it out of the UPS.

Figure 172. Pulling Out the Battery Pack

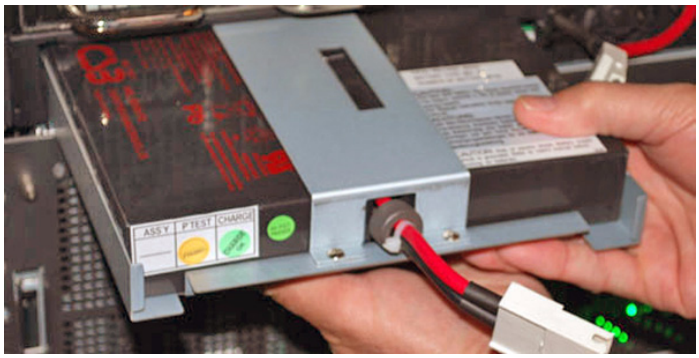

5. Slide the new battery pack into the UPS.
6. Reinstall the battery pack cover.
7. Connect the battery pack connector.
8. Reinstall the UPS bezel.
9. If you shut down the controller to change the battery pack, restart the controller now.

5.5.3.1 Testing Batteries

NOTE : It is recommended that the battery pack be charged for 6 to 24 hours before testing.

Press and hold the button for 3 seconds to initiate a self-test.

The 15-second test automatically distributes the load to the batteries and tests the battery pack's performance. While the test is in progress, the indicators cycle through and the alarm sounds. Upon completion, the UPS returns to normal mode as indicated by the and Load indicators.

If there is a problem with the battery pack, the alarm beeps, the indicator illuminates, and the indicator flashes. Check the battery pack connections and be sure the battery pack is fully charged. Call your service representative if the problem persists.

5.6 Controller Component Failures

The controller's Enclosure Fault LED turns amber when a fan failure, a power failure, or over temperature condition occurs (Figure 173).

Figure 173. Controller Enclosure Fault LED Indicator

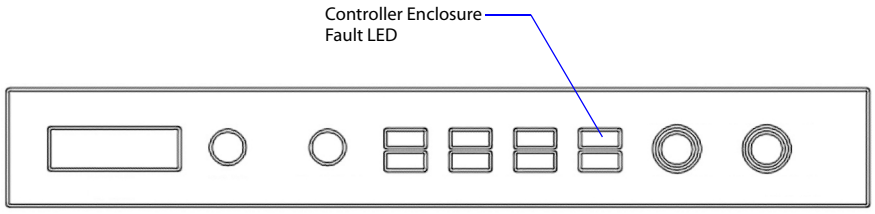

A single component failure, therefore, will not shut down the system. However, in the unlikely event of component failure, you can replace the failed component while the SFA10000 is running. The replaced component will automatically be returned to service once the component has been installed and booted up.

Warning

If the controller is powered up and you remove any module, replace it immediately. If the controller is used with modules missing for more than a few minutes, the controller can overheat, causing power failure and data loss. Such use will invalidate the warranty.

Warning

Observe all conventional ESD precautions when handling the controller components. Avoid contact with backplane components and module connectors.

5.6.1 Replacing a Power Supply Module

If a controller power supply fails, the Enclosure Fault LED on the front panel will turn **amber** and the green LED on the failed power supply will be **off**.

You can also determine if a power supply has failed using the CLUI command, **`SHOW POWER`**. The position of the failed power supply should match the label (PSU1 or PSU2) on the back of the controller.

If the power supply has failed, you must replace it. Once you remove it, you must replace it within 5 minutes to prevent the system from over-heating.

> **NOTE :** Obtain a replacement module of the same type before removing any faulty module.

Follow these steps to replace a power supply module:

1. Identify the failed power supply.
2. Disconnect its power cord.
3. Loosen the module's thumbscrew, and then slide the module out of the bay (Figure 174).

Figure 174. Controller Power Supply Module Removal

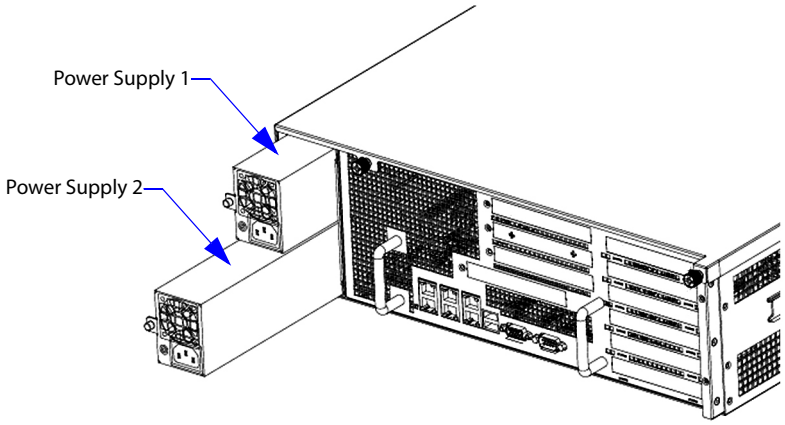

4. Slide the replacement module into the bay, making sure that it is fully inserted.
5. Tighten the thumbscrew to secure it.
6. Connect the power cord.
7. Verify that the Status LED is **green**, indicating that the module is operating normally.
8. Verify that the Enclosure Fault LED on the front panel is no longer amber.

5.6.2 Replacing a Fan Module

If there is any problem with the fan module, the Enclosure Fault LED on the front panel will turn **amber.**

You can determine if a fan has failed using the CLUI command, **`SHOW FAN`**. The position of the failed fan should match the label (FAN1, FAN2, FAN3 or FAN4) on the front of the controller.

If a fan has failed, you must replace it. Once you remove it, you must replace it within 5 minutes to prevent the system from over-heating.

NOTE : Obtain a replacement module of the same type before removing any faulty module.

Follow these steps to replace a fan module:

NOTE : Attempting to replace more than one fan simultaneously will cause the controller to fail.

1. Identify the failed fan.
2. Loosen the module's thumbscrew (Figure 175).
3. Slide the module out of the bay.

Figure 175. Controller Fan Removal

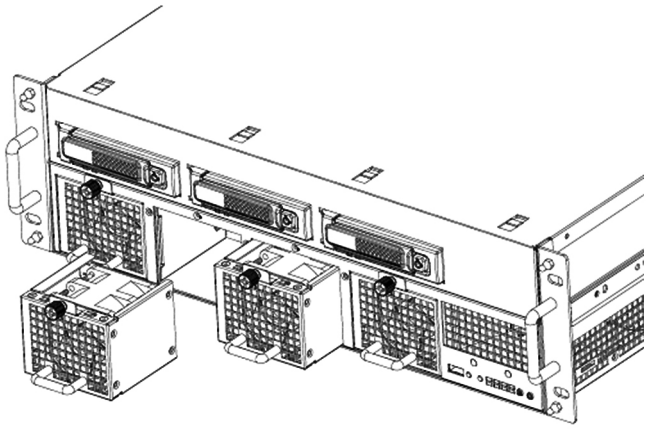

4. Slide the new module into the bay, making sure that it is fully inserted.
5. Tighten the thumbscrew to secure it.
6. Verify that the Enclosure Fault LED is no longer amber, indicating that the new module is operating normally.

5.6.3 Replacing an Internal Disk Module

Follow these steps to replace an internal disk module at the front of the controller:

1. Disengage the lock using the supplied key.
2. Press the handle release button to release the handle (Figure 176).
3. Pull the module out of the bay.

Figure 176. Controller Internal Disk Module Replacement

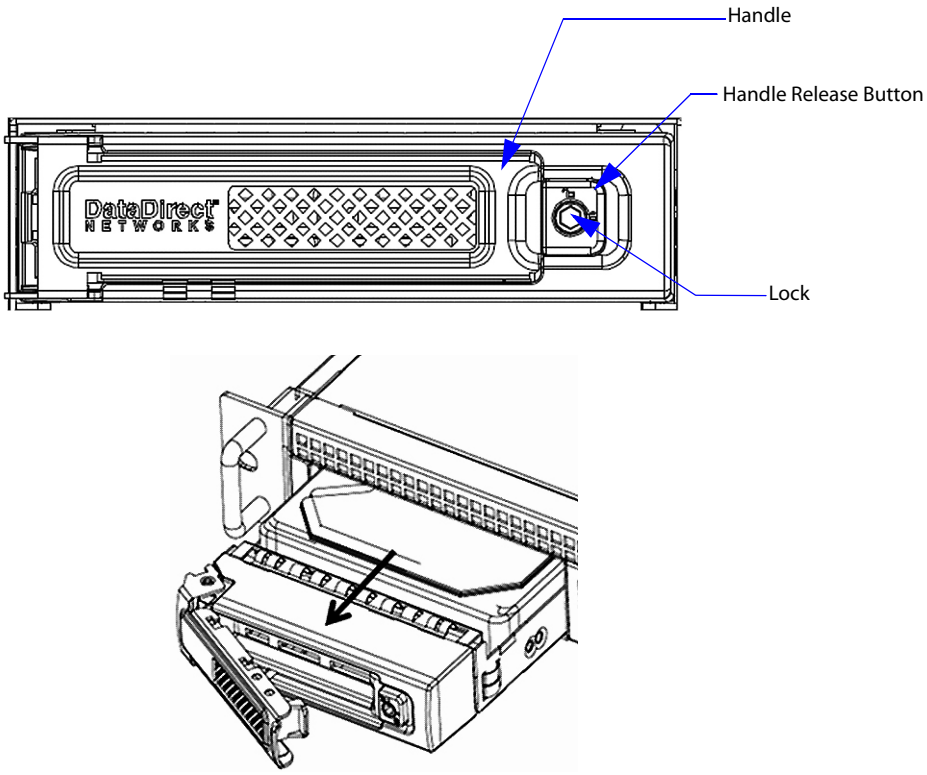

4. Slide the new module into the bay, making sure that it is fully inserted.
5. Close the handle and push the handle until it clicks, indicating that the handle is latched.

5.7 Disk Enclosure Component Failures

5.7.1 I/O Module Replacement

Please refer to the respective user guides if you need detailed instructions on how to physically replace an I/O module on the disk enclosure.

The following procedure describes how to replace (reseat or swap) a failing I/O module:

1. Identify the I/O module (IOM) that needs to be replaced.

   The need for an IOM replacement may have come from an event in an event log, the results of a **SHOW** command, an SNMP trap, or an email message. The event, SNMP trap, and email message include the element index within its enclosure and the enclosure WWN; the outputs of **SHOW ENCLOSURE** and **SHOW EXPANDER** have the same information. Use these and the **SHOW ENCLOSURE** command to determine the OIDEX of the IOM's enclosure.

2. Physically locate the IOM that needs to be replaced.

   First, find the ENCLOSURE and then find the EXPANDER within the enclosure. If the enclosures have been labeled with their WWNs then find the enclosure whose label matches the results from Step 1 above, otherwise use the **LOCATE ENCLOSURE** command with the enclosure OIDEX from Step 1. Viewing from the back of the enclosure, IOM 1 is on the left and IOM 2 is on the right. While you might be able to issue a **LOCATE EXPANDER** command using the enclosure OIDEX from Step 1, it requires that the IOM still be working well enough to light its Identify LED.

3. Identify the physical controller associated with the failing IOM.

   Follow the cable connected to the IOM that needs to be replaced back to the associated SFA controller. Each of the cables is labeled with a 3-digit number that has the form 0xx or 1xx. All of the cables with labels beginning with 0 go to one controller and all of those beginning with 1 go to the other controller. In cases where more than one SFA10000 is installed in a set of racks, each couplet reuses the same cable labels, so be careful not to confuse couplets.

4. Determine the index of the associated controller.

   Use the **LOCATE CONTROLLER <idx>** command to flash the blue Locate LED on the front panel.

5. Shut down the associated controller.

   Enter the **SHUTDOWN CONTROLLER <idx>** command where <idx> is index of the controller which was determined in Step 4. This will cause all pools to fail over to the other controller.

6. Physically replace the failing IOM.

7. Start the controller that was shut down.

   Press the Power button on the controller's UPS. Allow the controller to start. Eventually, pools will fail back to the controller that was started.

5.7.2 DEM Replacement on SS6000

Please refer to the *DDN SS6000 User Guide* if you need detailed instructions on how to physically replace a DEM.

The following procedure describes how to replace (reseat or swap) a failing DEM:

1. Identify the DEM that needs to be replaced.

   The need for a DEM replacement may have come from an event in an event log, the results of a **SHOW** command, an SNMP trap, or an email message. The event, SNMP trap, and email message include the element index within its enclosure and the enclosure WWN; the outputs of **SHOW ENCLOSURE** and **SHOW EXPANDER** have the same information. Use these and the **SHOW ENCLOSURE** command to find the OIDEX of the DEM's enclosure. Note that the DEMs in each SS6000 enclosure have expander indexes 3 through 10; indexes 1 and 2 are the I/O modules (IOM).

2. Physically locate the DEM that needs to be replaced.

   First, find the ENCLOSURE and then find the EXPANDER within the enclosure. If the enclosures have been labeled with their WWNs then find the enclosure whose label matches the results from Step 1, otherwise use the **LOCATE ENCLOSURE** command with the enclosure OIDEX from Step 1. While you might be able to issue a **LOCATE EXPANDER** using the enclosure OIDEX from Step 1, it requires that the DEM still be working well enough to light its Locate LED. Alternatively, **SHOW EXPANDER** provides the expander location and that corresponds to the label inside the top door of the SS6000 enclosure.

3. Identify the IOM that is associated with the failing DEM.

   SHOW EXPANDER gives a location that starts with DEM and ends with either A or B. Those DEMs with locations ending in A are associated with IOM 1 and those ending in B are associated with IOM 2.

4. Identify the physical controller associated with this IOM.

   Follow the cable connected to this IOM back to the associated SFA controller. Each of the cables is labeled with a 3-digit number that has the form 0xx or 1xx. All of the cables with labels beginning with 0 go to one controller and all of those beginning with 1 go to the other controller. In cases where more than one SFA10000 is installed in a set of racks, each couplet reuses the same cable labels, so be careful not to confuse couplets.

5. Determine the index of the associated controller.

   Use the **LOCATE CONTROLLER <idx>** command to flash the blue Locate LED on the front panel.

6. Shut down the associated controller.

   Enter the **SHUTDOWN CONTROLLER <idx>** command where <idx> is index of the controller which was determined in Step 5. This will cause all pools to fail over to the other controller.

7. Physically replace the failing DEM.

8. Start the controller that was shut down.

   Press the Power button on the controller's UPS. Allow the controller to start. Eventually, pools will fail back to the controller that was started.

5.7.3 SAS Cable Replacement

The following procedure describes how to replace (reseat or swap) a failing cable:

1. Physically locate the cable that needs to be replaced.

   This is usually done by physically inspecting the cables or by noting that one of the green SAS link LEDs on an I/O module (IOM) is not lit.

2. Identify the IOM that is associated with the failing cable.

   This is done by physically following the cable to an IOM.

3. Identify the physical controller associated with this IOM.

   Follow the cable connected to this IOM back to the associated SFA controller. Each of the cables is labeled with a 3-digit number that has the form 0xx or 1xx. All of the cables with labels beginning with 0 go to one controller and all of those beginning with 1 go to the other controller. In cases where more than one SFA10000 is installed in a set of racks, each couplet reuses the same cable labels, so be careful not to confuse couplets.

4. Determine the index of the associated controller.

   Use the **`LOCATE CONTROLLER <idx>`** command to flash the blue Locate LED on the front panel.

5. Shut down the associated controller.

   Enter the **`SHUTDOWN CONTROLLER <idx>`** command where <idx> is index of the controller which was determined in Step 4. This will cause all pools to fail over to the other controller.

6. Physically replace the failing cable.

7. Start the controller that was shut down.

   Press the Power button on the controller's UPS. Allow the controller to start. Eventually, pools will fail back to the controller that was started.

Appendices

A

Technical Specifications

Here are the technical specifications for the SFA10000/10000E.

NOTE : Specifications subject to change without notice.

Configuration, Performance, & Capacity

Host side technology (SFA10000)	(16) 8Gb Fibre Channel or (8) 40Gb InfiniBand
Host side technology (SFA10000E)	(16) individual InfiniBand or 10G Ethernet
Disk side technology	(40) 4 × 3Gb/s SAS
Supported disk technology	SAS / SATA / SSD
Management interface	RS-232 and Ethernet (SSH)
Cache memory (SFA10000)	24GB of RAM, 12GB ECC/RAID protected cache
Cache memory (SFA10000E)	24GB ECC/RAID protected cache
Application memory (SFA10000E)	96GB ECC protected cache
Number of disk modules supported	Up to 1200 disks, 2.4 Petabytes per array
Number of storage pools supported	256
Number of virtual disks supported	512; 64 per storage pool
Number of Spare Pools supported	16
Hot spare capability	Yes
RAID parity protection	1+1 RAID 1, 4+1 or 8+1 RAID 5, 4+2 or 8+2 RAID 6
Throughput (SFA10000)	12GB/s sustained with large well aligned, sequential IO
Throughput (SFA10000E)	6GB/s sustained
IOPS (SFA10000)	1M to cache; 300,000 to Disk
IOPS (SFA10000E)	500,000 to cache; 150,000 to Disk

Reliability

SES (SCSI Enclosure Services) protocol support	Yes
Temperature monitoring	Yes
Battery-backed write-back cache	Yes
Redundant hot-swappable power modules	Dual-redundant

Reliability

Redundant hot-swappable cooling modules	N+1
Redundant controllers	Dual-redundant
Redundant disk enclosures	N+1

Physical, Power & Environmental

Active/Active Dual Controllers	Dimensions	Height: 14" (356 mm) (includes 2 UPS units) Width: 17" (432 mm) Depth: 25.5" (648 mm)
	Weight	120 lbs (54.5 kg) Controllers only; 248 lbs (112.5 kg) with UPS units
	Voltage range	200-240 VAC @ 47-63 Hz
	Average power	1200 W
	Average cooling	4095 BTU/hr
SS6000 Disk Enclosure	Dimensions	Height: 7" (178 mm) Width: 17" (432 mm) Depth: 36" (914 mm) without bezel; 42" (1067 mm) with cable management arms
	Weight	240 lbs (109 kg) with disk modules; 120 lbs (54.5 kg) without disk modules
	Voltage range	200-240 VAC @ 47-63 Hz
	Average power	1750 W
	Average cooling	5973 BTU/hr
SS7000 Disk Enclosure	Dimensions	Height: 6.97" (177 mm) Width: 16.56" (420.6 mm) Depth: 34" (863.6 mm) without bezel and cable management arms; 38" (965.2 mm) with bezel and cable management arms
	Weight	215 lbs (97.7 kg) with disk modules; 105 lbs (47.7 kg) without disk modules
	Voltage range	190-264 VAC @ 47-63 Hz
	Power Supply	1865 W

Physical, Power & Environmental

SS2460 Disk Enclosure	Dimensions	Height: 3.5" (890 mm) Width: 17.6" (447 mm) Depth: 19.3" (490 mm)
	Weight	53 lbs (27 kg) with disk modules
	Voltage range	100-240 VAC @ 50-60 Hz
	Maximum average output power	550 W

Operating environment (temperature / relative humidity)	5 °C to 35 °C / 20% - 80%, non-condensing
Non-operating environment (temperature/ relative humidity)	-10 °C to 50 °C / 20% - 80%, non-condensing
Agency certifications	UL, CE, cUL, C-Tick, FCC

NOTE : Specifications subject to change without notice.

B

Disk Module Installation

This appendix provides detailed instructions on how to install a disk module in the SS6000, SS7000, and SS2460 enclosures.

B.1 SS6000 Disk Module Installation Procedure

NOTE : The SS6000 design allows for disk bays to be left empty without the need for fitting dummy disk modules.

1. Squeeze the release tabs together on the front panel and carefully pull the enclosure forward out of the rack. Keep pulling until the enclosure locks and you hear a clicking sound.
2. Disengage both enclosure cover latches (Figure 177) and open the covers.

Figure 177. SS6000 Enclosure Cover Latch

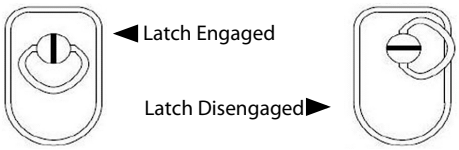

3. On the disk module, slide the latch backward to release the handle (Figure 178).

Figure 178. Release Disk Module Handle

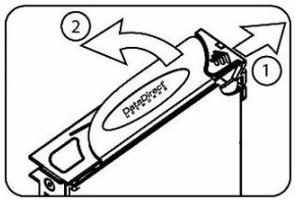

1. Slide Latch Backward
2. Lift Handle

4. Insert the module into a disk bay. Cam the disk module home. The camming foot on the base of the module will engage into the slot in the enclosure.
5. When the module is fully inserted, close the handle. You should hear a click as the latch engages and holds the handle closed.
6. After you have installed all the disk modules in this enclosure, close the enclosure covers and engage both cover latches.
7. Squeeze the tabs on both rack slides and push the enclosure back into the rack.

B.2 SS7000 Disk Module Installation Procedure

NOTE : The SS7000 design allows for disk bays to be left empty without the need for fitting dummy disk modules.

1. Squeeze the release tabs together on the front panel (Figure 179) and carefully pull it forward out of the rack. Keep pulling until the enclosure locks and you hear a clicking sound.

Figure 179. Release Tabs on Front Panel

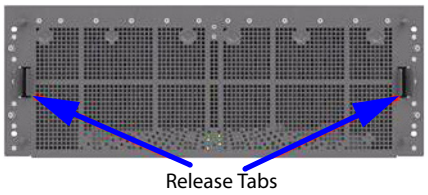

2. Open the covers (Figure 180):
   - Loosen the two Phillips screws on the covers
   - Press the release buttons on the top cover release handles
   - Pull up the release handles

Figure 180. Open Top Covers

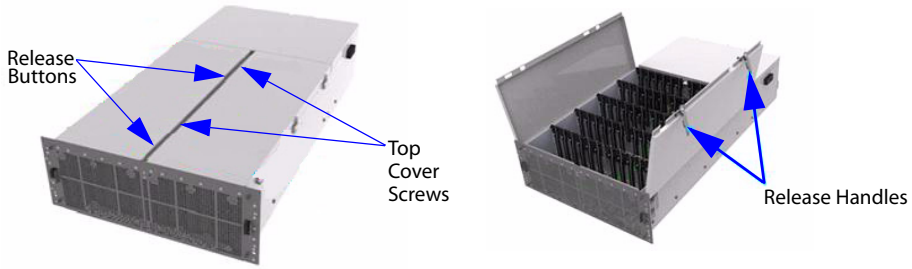

3. On the disk module, press the release catch to release the handle (Figure 181). Carefully slide the module into a disk slot until it contacts the baseboard and begins to engage the handle. Press down on the handle until it latches which will cam the module into place.

Figure 181. Insert Disk Module

4. After you have installed all the disk modules in this enclosure, close the enclosure covers and engage both cover release handles. Replace the two Phillips screws.

5. Squeeze the tabs on both rack slides and push the enclosure back into the rack.

B.3 SS2460 Disk Module Installation Procedure

NOTE : To maintain proper airflow and cooling, a disk module or a dummy disk module must be installed in every slot of the enclosure.

1. Press the handle release button on the disk module to release the handle (Figure 182).

Figure 182. Disk Module Handle Release Button

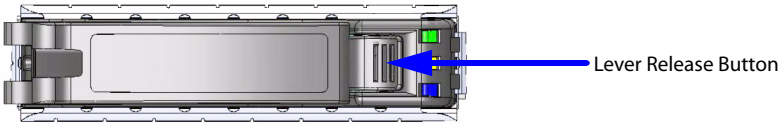

2. Hold the disk module so that the LEDs are on the bottom and the handle is fully open.
3. Slide the disk module into the slot until the handle starts to close (Figure 183).

Figure 183. Insert Disk Module into Disk Slot

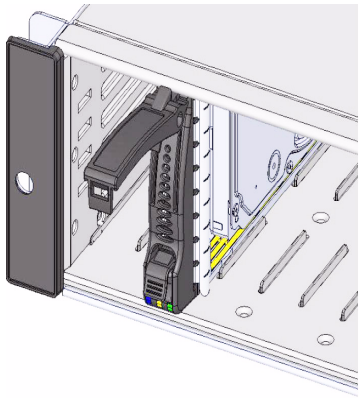

4. Gently push the handle down to cam the module home. You should hear a click as the latch engages and holds the handle closed (Figure 184).

Figure 184. Close Handle on Disk Module

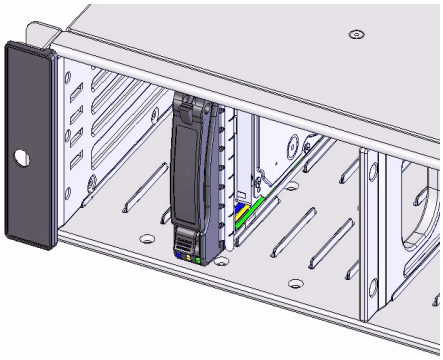

C

SS6000 Disk Maps for SFA10000E

The charts below illustrate the disk maps for proper placement in the SS6000 disk enclosures when used with the SFA10000E.

Disk Enclosure 1 (Top)											
49	50	51	52	53	54	55	56	57	58	59	60
37	38	39	40	41	42	43	44	45	46	47	48
25	26	27	28	29	30	31	32	33	34	35	36
13	14	15	16	17	18	19	20	21	22	23	24
1	2	3	4	5	6	7	8	9	10	11	12

▲ Front of Rack ▲

Disk Enclosure 6											
349	350	351	352	353	351	355	356	357	358	359	360
337	338	339	340	341	342	343	344	345	346	347	348
325	326	327	328	329	330	331	332	333	334	335	336
313	314	315	316	317	318	319	320	321	322	323	324
301	302	303	304	305	306	307	308	309	310	311	312

▲ Front of Rack ▲

Disk Enclosure 2											
109	110	111	112	113	114	115	116	117	118	119	120
97	98	99	100	101	102	103	104	105	106	107	108
85	86	87	88	89	90	91	92	93	94	95	96
73	74	75	76	77	78	79	80	81	82	83	84
61	62	63	64	65	66	67	68	69	70	71	72

▲ Front of Rack ▲

Disk Enclosure 7											
409	410	411	412	413	414	415	416	417	418	419	420
397	398	399	400	401	402	403	404	405	406	407	408
385	386	387	388	389	390	391	392	393	394	395	396
373	374	375	376	377	378	379	380	381	382	383	384
361	362	363	364	365	366	367	368	369	370	371	372

▲ Front of Rack ▲

Disk Enclosure 3											
169	170	171	172	173	174	175	176	177	178	179	180
157	158	159	160	161	162	163	164	165	166	167	168
145	146	147	148	149	150	151	152	153	154	155	156
133	134	135	136	137	138	139	140	141	142	143	144
121	122	123	124	125	126	127	128	129	130	131	132

▲ Front of Rack ▲

Disk Enclosure 8											
469	470	471	472	473	474	475	476	477	478	479	480
457	458	459	460	461	462	463	464	462	466	467	468
445	446	447	448	449	450	451	452	453	454	455	456
433	434	435	436	437	438	439	440	441	442	443	444
421	422	423	424	425	426	427	428	429	430	431	432

▲ Front of Rack ▲

Disk Enclosure 4											
229	230	231	232	233	234	235	236	237	238	239	240
217	218	219	220	221	222	223	224	225	226	227	228
205	206	207	208	209	210	211	212	213	214	215	216
193	194	195	196	197	198	199	200	201	202	203	204
181	182	183	184	185	186	187	188	189	190	191	192

▲ Front of Rack ▲

Disk Enclosure 9											
529	530	531	532	533	534	535	536	537	538	539	540
517	518	519	520	521	522	523	524	525	526	527	528
505	506	507	508	509	510	511	512	513	514	515	516
493	494	495	496	497	498	499	500	501	502	503	504
481	482	483	484	485	486	487	488	489	490	491	492

▲ Front of Rack ▲

Disk Enclosure 5											
289	290	291	292	293	294	295	296	297	298	299	300
277	278	279	280	281	282	283	284	285	286	287	288
265	266	267	268	269	270	271	272	273	274	275	276
253	254	255	256	257	258	259	260	261	262	263	264
241	242	246	244	245	246	247	248	249	250	251	252

▲ Front of Rack ▲

Disk Enclosure 10											
589	590	591	592	593	594	595	596	597	598	599	600
577	578	579	580	581	582	583	584	585	586	587	588
565	566	567	568	569	570	571	572	573	574	575	576
553	554	555	556	557	558	559	560	561	562	563	564
541	542	543	544	545	546	547	548	549	550	551	552

▲ Front of Rack ▲

D

SS6000 Zoning Configuration Procedure

This appendix provides information on how to change the zoning configuration on an SS6000 enclosure.

The disk enclosures must be zoned correctly. The zoning configuration is determined by the piano switches located behind the front bezel of the enclosure. Follow these steps to verify the zoning configuration on each disk enclosure:

1. Slide the enclosure out from the rack by squeezing the tabs on both rack slides. Keep pulling until the enclosure locks and you hear a clicking sound.
2. Remove the front bezel by removing the four screws (Figure 185).

Figure 185. Removing SS6000 Front Bezel

3. Once the front bezel is removed, it is possible to see the piano switches through the grey bezel (Figure 186). If the switches are hidden by a Mylar switch cover, remove the cover with a small pair of tweezers. If the switches are still not visible, remove the grey bezel. It is attached to the chassis with 32 small screws (18 on the front and 14 on the bottom).

Figure 186. Zoning Configuration Piano Switches Behind Bezel

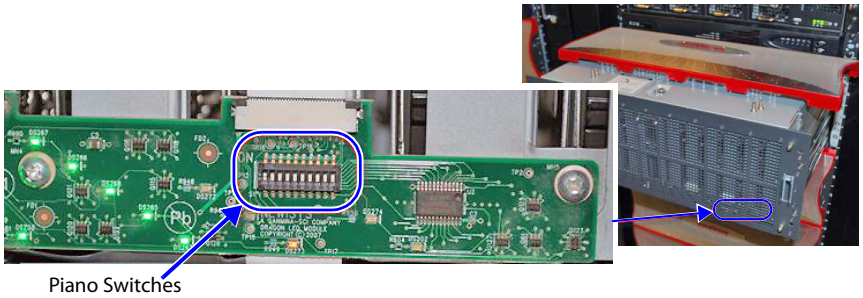

4. Verify that the piano switches are correctly set.

   For system with 5 disk enclosures (zoning configuration 1), the left-most piano switch should be up and the rest down. This divides the disk slots into two logical enclosures.

   For system with 10 or 20 disk enclosures (zoning configuration 0), all the piano switches should be down. This connects all disk slots into one logical enclosure.

   If a change is needed, remove the Mylar switch cover with a small pair of tweezers and then use a dental tool to move the switches. The tweezers and dental tool can be

inserted through the grey bezel air holes. If changing the switches with a dental tool proves to be too difficult, then the grey bezel can be removed.

5. Replace the front bezel (and the grey bezel, if previously removed).
6. Push the enclosure back into the rack.
7. Repeat the above steps on all disk enclosures.

E

Rack Installation Procedure

This appendix provides information on how to mount a SFA10000 controller and UPS unit in a rack.

E.1 Rackmounting the Controller

The SFA10000 controller is designed to fit within a 3U rack space. Follow these steps to mount the controller in your rack:

1. Figure 187 shows the rackmount kit for the controller. There are two sets of rail and two bags of screws included.

Figure 187. Controller Rackmount Kit

2. The two mounting brackets should already be attached to the controller. If not, attach the brackets to the two sides of the controllers using six screws (Figure 188).

Figure 188. Mounting Bracket on Controller

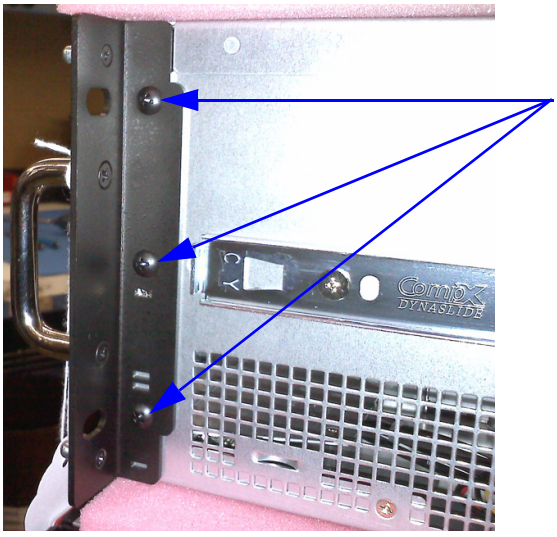

3. Pull on the “Front-Release” to unlock the inner rail from the slide assembly (Figure 189). Then release the Detent-Lock and push the inner rail inwards to retract the inner rail.

Figure 189. Slide Assembly

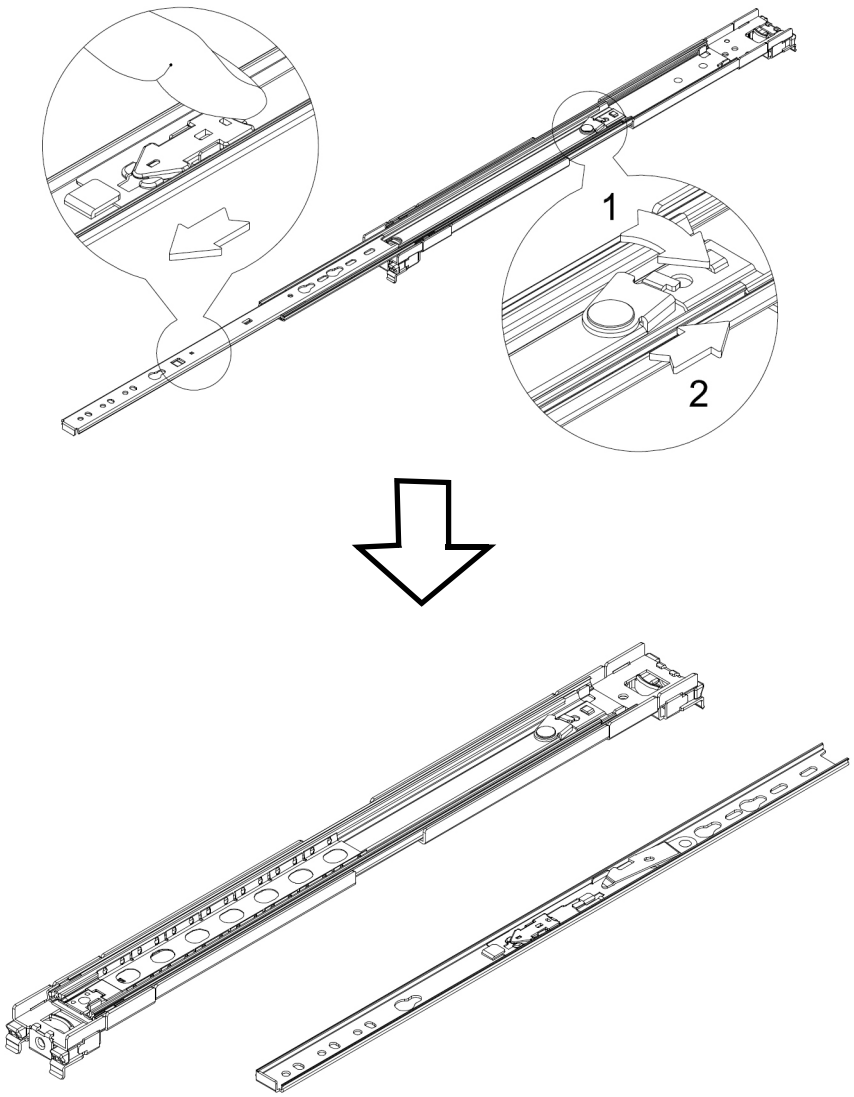

4. Align the front bracket with the mounting holes in rack (Figure 190). Optional: Remove metal spacer for aluminum racks.

Figure 190. Attach Slides to Rack (1)

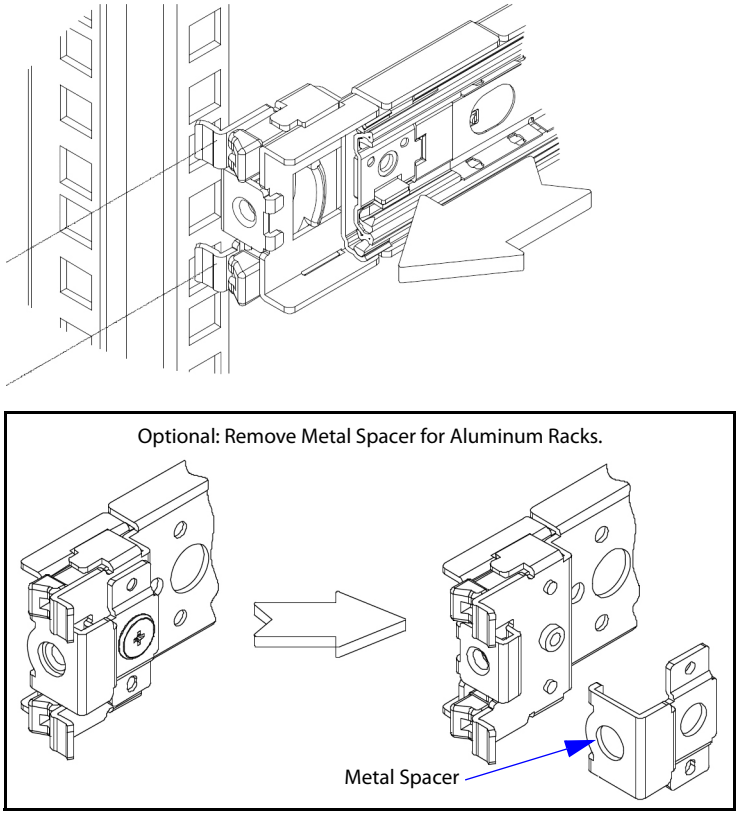

5. Push the front bracket onto rack (Figure 191). Optional: M6x10L screws can be used to secure the rails with posts if needed.

Figure 191. Attach Slides to Rack (2)

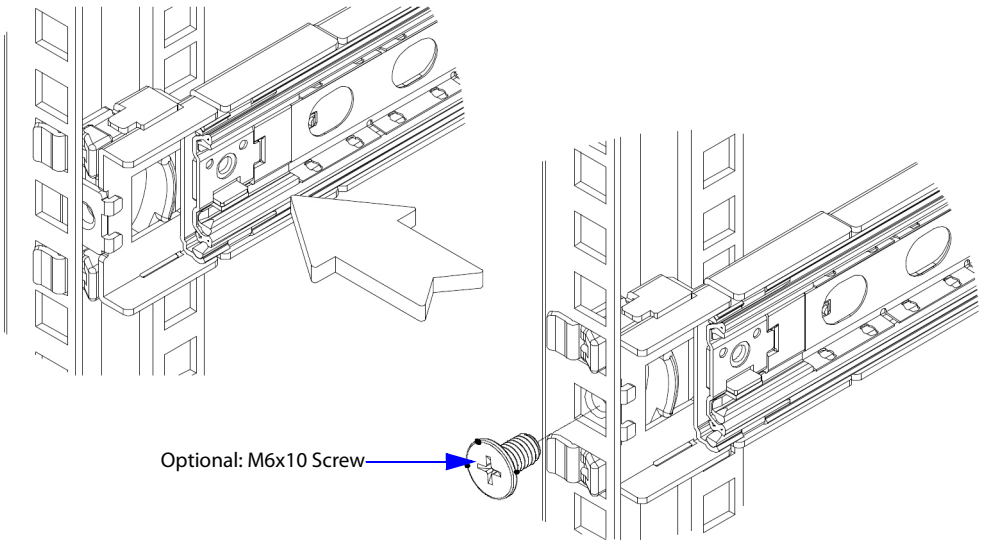

6. Similarly, attach the end bracket to the rear of rack (Figure 192).

Figure 192. Attach Slides to Rack (3)

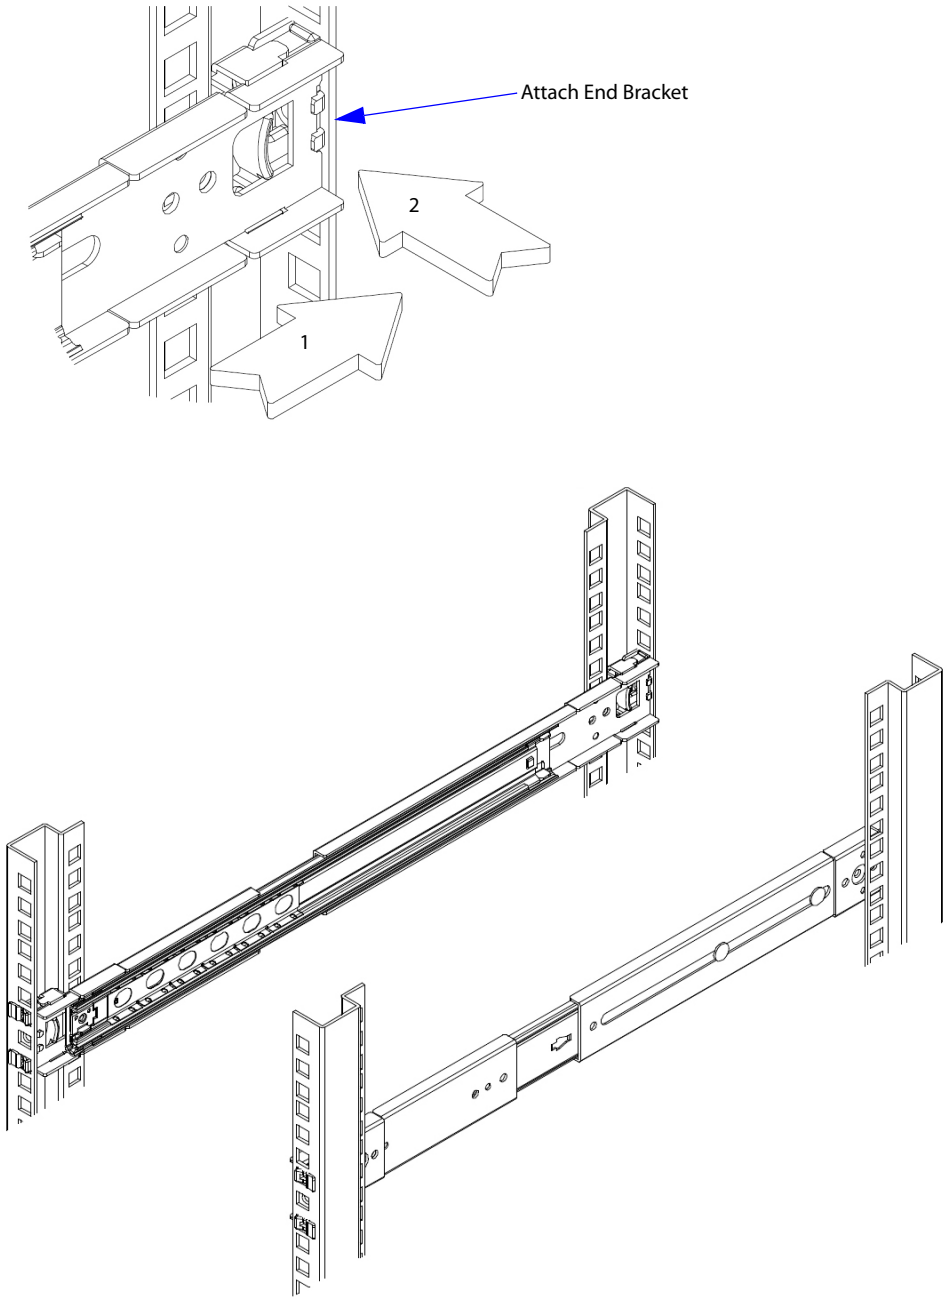

7. Slide the two inner rails onto the chassis hooks on both sides of the controller (Figure 193).

Figure 193. Attach Inner Rails to Chassis

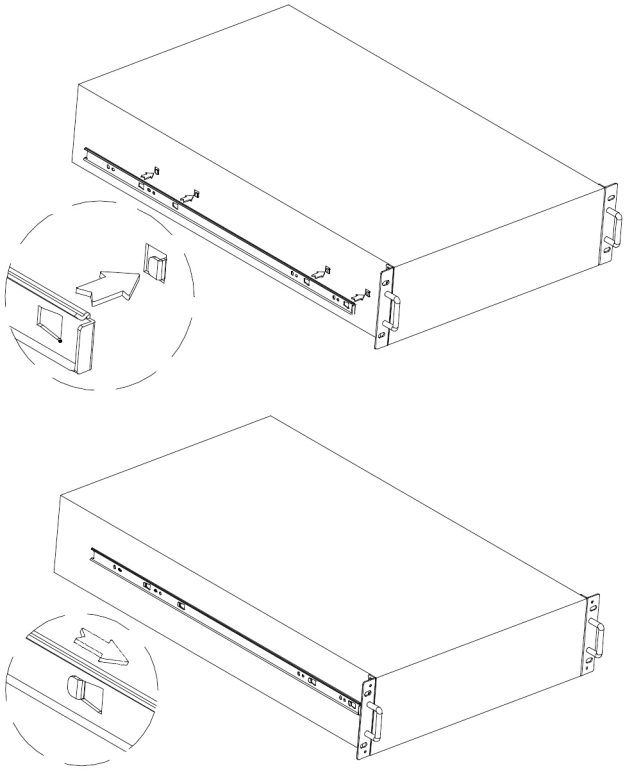

8. Secure the inner rails using two screws on each side near the rear of chassis (Figure 194).

Figure 194. Secure Inner Rails

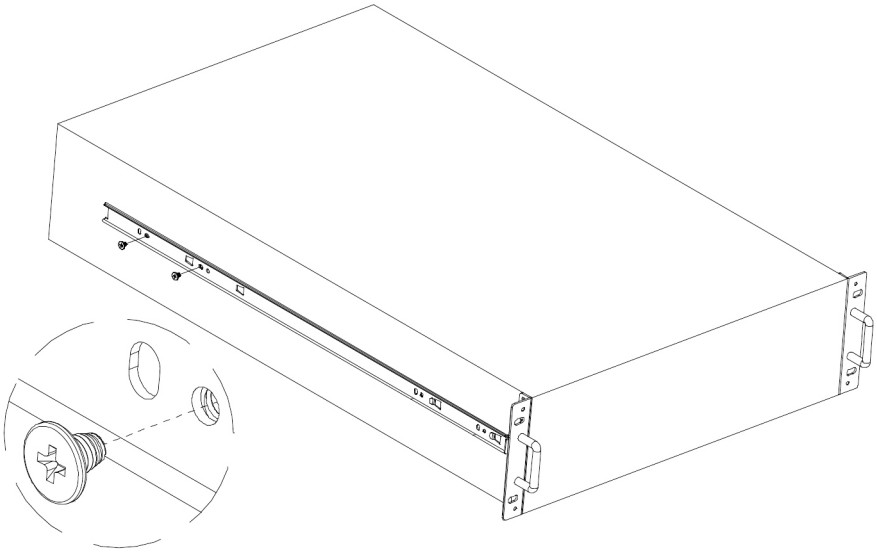

9. Then push the chassis into slides to complete rack installation (Figure 195).

CAUTION ! The main purpose of the chassis handles is for rackmount use only. The controller is NOT to be carried or transported by the handles.

Figure 195. Install Chassis to Rack

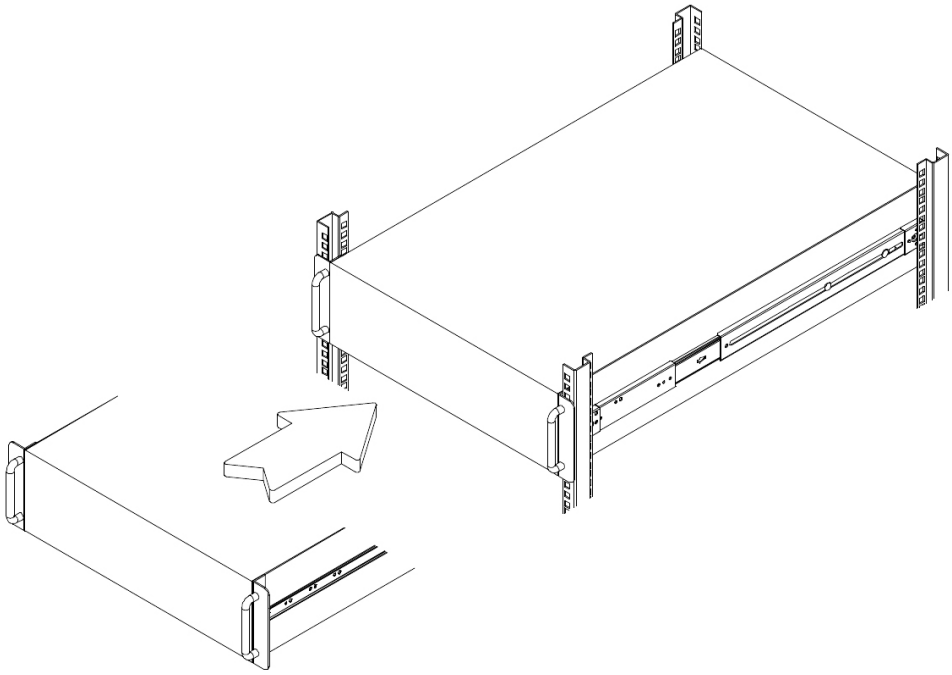

E.2 Rackmounting the UPS Unit

The UPS unit is designed to fit within an 1U rack space. Follow these steps to mount the UPS in your rack:

1. The rackmount kit for the UPS includes two mounting brackets with 8 screws and two rack rails with 8 screws. Figure 196 shows a rackmount rail.

Figure 196. UPS Rackmount Rail

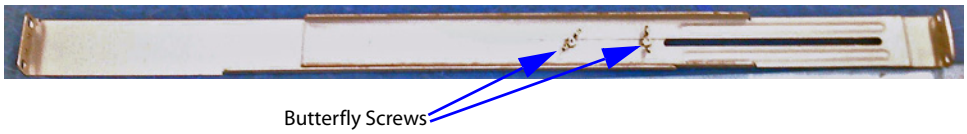

2. Attach the two mounting brackets to the two sides of the UPS using 8 flat head screws (Figure 197). Use the bracket marked “LEFT” for the left side, when viewing from front of unit, and the “RIGHT” bracket for the right side.

Figure 197. Mounting Bracket on UPS

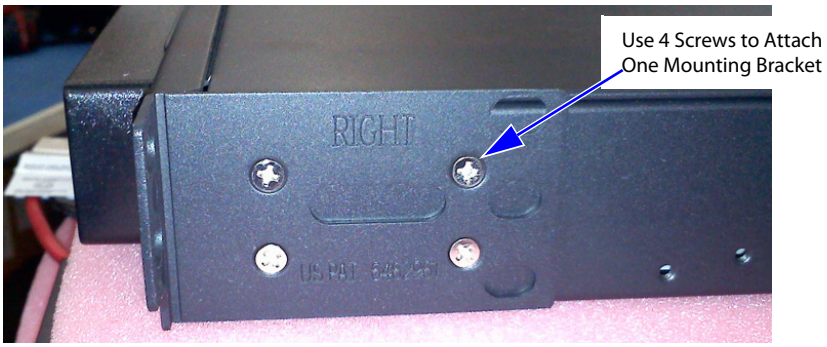

3. Loosen the two butterfly screws on each rack rail (Figure 196). Adjust the length of the rails to fit the depth of your rack. Then tighten the butterfly screws.

4. Use one screw at the bottom to secure the front of the rail to the rack (Figure 198).

Figure 198. Secure Front of Rail to Rack

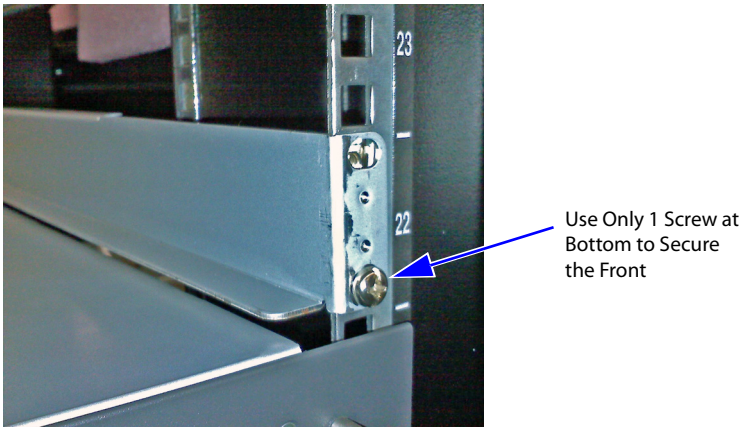

5. Use two screws to secure the rear end of the rail to the rack (Figure 199).

Figure 199. Secure Rear End of Rail to Rack

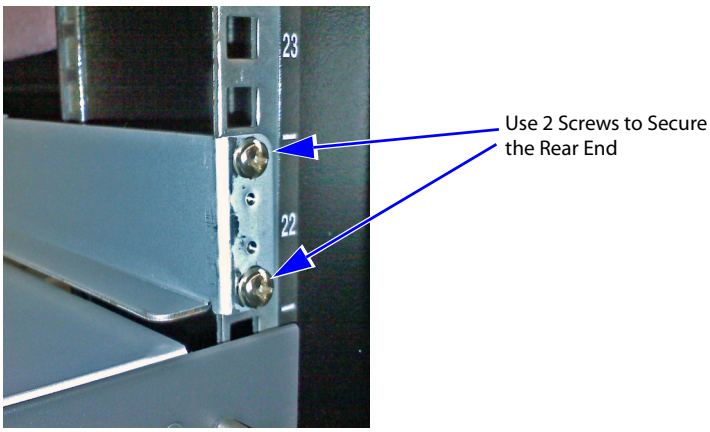

6. Slide the UPS onto the rack rails. Then use one screw on each side to secure the mounting brackets to the rack (Figure 200).

Figure 200. Secure UPS to Rack

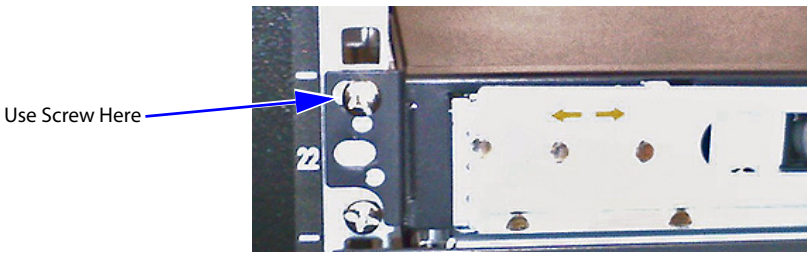

F

FRU Part Numbers

This appendix provides the list of field replaceable unit (FRU) part numbers.

Controller FRU

FRU Part Number	Unit per Controller	Description
FRU-10K-PS	2	Power supply module
FRU-10K-FM	4	Fan module
FRU-10K-RLK	1	24-inch rail kit
FRU-10K-IHD	3	1TB SATA Internal hard drive (cache protect and firmware)
FRU-10K-BBU	1	UPS (includes UPS, rail kit, and cables)
FRU-10K-BBUB	1	UPS (FRU-10K-BBU) replacement battery
10KS-BLKFC-FR		SFA10000 controller with 8 FC-8 ports
10KS-BLKIB-FR		SFA10000 controller with 4 QDR IB SRP ports

SS6000 Disk Enclosure FRU

FRU Part Number	Unit per Enclosure	Description
FRU-60S-RLK3	1	Adjustable rail kit
FRU-60S-CMA2	1	Cable management arms
FRU-60S-PS3	2	Power supply module
FRU-60S-DEM	8	DEM module
FRU-60S-IOM2	2	I/O module
KIT-CBL10KSS		SAS cable kit to connect controller to disk enclosures (includes 8 x 2-meter cables and 12 x 3-meter cables)
CBL-MSMS-1M		1-meter SAS cable to connect controller to disk enclosure
CBL-MSMS-2M		2-meter SAS cable to connect controller to disk enclosure
CBL-MSMS-3M		3-meter SAS cable to connect controller to disk enclosure
CBL-MSMS-4M		4-meter SAS cable to connect controller to disk enclosure
CBL-MSMS-5M		5-meter SAS cable to connect controller to disk enclosure

SS6000 Disk Enclosure FRU

FRU Part Number	Unit per Enclosure	Description
CBL-MSMS-6M		6-meter SAS cable to connect controller to disk enclosure
601-BLANK		Empty slot filler for enclosure

SS7000 Disk Enclosure FRU

FRU Part Number	Unit per Enclosure	Description
FRU-SS7K-RK	1	Adjustable rail kit
FRU-SS7K-CMA	1	Cable management arms
FRU-SS7K-PS	2	1865 Watt Power supply module
FRU-SS7K-IO	2	I/O module
FRU-SS7K-KIT		Spares kit includes 1 x 1865W power supply module and 1 x I/O module
KIT-CBL10KSS		Mini-SAS cable kit to connect an SFA10000 singlet or a SFA10000E couplet to 5 x 60-bay enclosures. Also used to connect SFA10000 couplet to 5 x 24-bay enclosures (includes 2 x 2-meter cables and 18 x 3-meter cables)
CBL-MSMS-XM		Cable kit includes various length of mSAS/mSAS cables to connect SFA10000 singlet to 60-bay enclosures
SS7K-BLANK		Empty slot filler enclosure

SS2460 Disk Enclosure FRU

FRU Part Number	Unit per Enclosure	Description
FRU-2460-RLK	1	Adjustable rail kit
FRU-2460-PSU	2	Power supply module
FRU-2460-IOM	2	I/O module
KIT-CBL10KSS		Mini-SAS cable kit to connect an SFA10000 singlet or a SFA10000E couplet to 5 x 60-bay enclosures. Also used to connect SFA10000 couplet to 5 x 24-bay enclosures (includes 2 x 2-meter cables and 18 x 3-meter cables)
FRU-2460-ENC		Empty chassis
FRU-2460-BNK		2.5" dummy disk module

G

Terminology

AP	Application Processor
Application Stack	A layer of SFA OS that provides services external to the SFA appliance through the use of Virtual Disks. For example, Fibre Channel block device.
Client Channel	A port on a Client I/O Controller that can be used by hosts to communicate with an Application Stack.
Client I/O Controller	A physical interface (FC or IB) that has a unique PCI bus number and is utilized for one or more Client Channels. A Client I/O Controller is managed by a single Application Stack.
Couplet	Dual controller subsystem
Discovered Initiator	A volatile object that represents an FCP or SRP initiator port that is discovered on the fabric.
Disk Channel	A port on a Disk I/O Controller that can be connected to one or more Physical Disks and Disk Enclosures.
Disk I/O Controller	A physical interface (SAS) that has a unique PCI bus number and is utilized for one or more Disk Channels. A Disk I/O Controller is managed by the RAID Stack running on a single RAID Processor.
GiB (gibibyte)	A unit of digital information storage and denote 1,073,741,824 bytes (2^{30} bytes).
Host	A persistent logical object that represents an abstract client for one or more Virtual Disk's and that is given access to data on those Virtual Disk's.
ICL	Inter-Controller Link
Initiator	A persistent logical object associated with a Host that represents a FCP or SRP initiator port on a Host's HBA.
KiB (kibibyte)	A unit of digital information storage and denote 1,024 bytes (2^{10} bytes).
MiB (mebibyte)	A unit of digital information storage and denote 1,024 kibibytes (2^{20} bytes).

Physical Disk (PD)	A disk module (SAS or SATA) or solid-state device (SSD) used by the system back-end to store data and subsystem metadata.
Pool	A set of physical disks. There are three types: Storage Pool, Spare Pool, and Unassigned Pool.
Presentation	A persistent logical object that describes an association between a Host and a Virtual Disk that specifies if a Host may access the Virtual Disk and any specific constraints for access.
RAIDset	A container formed by a set of Physical Disk extents and controlled by one of the RAID algorithms, such as RAID6.
RP	RAID Processor
Singlet	Single Controller
Spare Pool	A set of physical disks that are used as spares in one or more storage pools.
Storage Pool	A set of physical disks whose extents may be parts of RAIDset's.
Unassigned Pool	The set of physical disks that are not assigned to any pool.
Virtual Disk (VD)	A set of disk blocks that may be presented to an operating system. It is also the primary logical object realized by the back-end.

Index

A

AC power requirement 17
airflow 16
AP 173
application processor, controller 9
application stack 173

B

background job 85
battery, replace UPS 142

C

cable, replace 152
cache
  mirrored write back, set 83
  write back, set 83
CLUI, using 51
controller
  application processor 9
  fan 8
  fan, replace 148
  hardware 5
  internal disk modules 9
  internal disk, replace 149
  IO ports 10
  LED, status 7
  power supply 8
  power supply, replace 147
  RAID processor 9
cooling requirements 17
create
  spare pool 44, 66
  storage pool 40
  virtual disk 42
customer service, contacting 178

D

DataDirect Networks, contacting 178
date, set system 78
delete
  storage pool 62
  virtual disk 64
DEM, replace SS6000 151
DIP switch, UPS 32
discovered initiators 46, 67
disk
  fail a 81
  failure recovery 139
  fail, clear 81
  information 55
  internal, controller 9
  listing all 54
  location, physical 56
  replace, manual 81
  SATA, power cycle 141
  states 55
disk enclosure
  cable connections 22
  DEM, replace SS6000 151
  I/O module, replace 150
  zoning 161
disk failure recovery 139
disk map, SFA10000E 160
disk module
  install 20, 157, 161, 163
disk timeout 80

E

email notification 102
Ethernet connection 101
event log 136
  display 99
  structure 100
  terms 100

F

failure
  disk, fixing 139
  fan, controller 148
  power supply, controller 147
fan
  controller 8
  controller, replace 148
features, system 3
field replaceable unit, list of 171
firmware
  current version 89
  update 88
  update controller 89
  update disk enclosure 92
  update physical disk 96
floor loading 17
fractional rebuild 80
FRU
  controller 171
  SS2460 172
  SS6000 171
  SS7000 172
full rebuild 80

G

GUI 107

H

host, connect controller to 29
host, presentation to 46, 67

I

ICL 21
import discovered initiator 70
initialization 60
    failed 60
installation 14
installation planning 15
inter-controller link 21
internal disk, replace controller 149
IO ports, controller 10
IP address, set controller 101
I/O module, replace 150

J

job priority, set 85

L

LED
    controller 7
    controller IO port 11
    internal disk module, controller 9
    UPS 12
loading, floor 17
login 52
logout 52, 102
log, event 99, 114

M

Manual Intervention Required 137
masking 49
MIR 137
    clear, state 137
    conditions 137
mirrored write back cache, set 83

N

name, set subsystem 37
network interface, change 101
Network Time Protocol 78, 110
NTP, set 78, 110

P

packing list 19
partial rebuild 80
persistent reservation, SCSI 69
physical location
    disk 56
    storage pool 62
    unassigned pool 56
pool
    spare, create 44, 66
    spare, display 65
    storage, create 40
    unassigned 54
power cycle disk drive 141
power requirement 17
power supply
    controller 8
    replace controller 147
presentation 46, 67
    MacOS 69
    masking 49
priority
    job, set 85
    rebuild, set 86

R

rack location 16
RAID processor, controller 9
rebuild
    automatic, set 80
    fractional 80
    full 80
    manual, initialize 81
    partial 80
    set priority 86
repair, how to return product for 178
replace disk, manual 81
restart
    controller 82
    subsystem 82
right side I/O 84
RMA, obtaining 178
RS-232 interface set up 51
RS-232 terminal set up 51

S

sales, contacting 178

SAS cable, replace 152
SATA disk, power cycle 141
SATAssure 61
SATAssure, set 82
SCSI persistent reservation 69
serial port, set up 51
service clearance 16
SFA 2
shutdown
  controller 82
  subsystem 82
site preparation 15
SNMP set up 103
spare pool
  create 44, 66
  delete 66
  display 65
  global 65
  name, change 66
  view list of 64
sparing policy 80
specifications, technical 154
SSH 101
SS6000 disk map 160
storage pool
  configuration, checking 56
  create 40, 59
  delete 62
  initialization 60
  list of, view 57
  location, physical 62
  name, change 62
  SATAssure, set 82
  sparing policy 80
  status 57
  verify 62

T

technical specifications 154
technical support, contacting 178
Telnet 101
terminology 173
time, set system 78

U

unassigned
  location, physical 56
unassigned pool 54
UPS 12
  battery pack, replace 142
  battery, connect 31
  connect to controller 31
  DIP switch 32
  test batteries 145

V

verify pool 62
version information 89
virtual disk
  assign to host 71
  associated with storage pool, view 59
  counters 84
  create 42, 64
  delete 64
  name, change 64
  status, checking 63
visual indication
  disk 56
  storage pool 62
  unassigned pool 56

W

wedged disk, troubleshoot 141
write back cache, set 83

Z

zoning configuration, disk enclosure 161

Contacting Technical Support & Shipping Instructions

You can contact DataDirect Networks **Technical Support** by:
Telephone at 1.888.634.2374 (North America)
1.818.718.8507 (International)
Internet email at support@ddn.com

Technical Support is available Monday through Friday.
After hours support is available to those customers that have a Service Contract. Please contact our sales representatives for more information.

You can find us on the World Wide Web at www.ddn.com

You can contact our **Customer Service** by:
Telephone at 1.818.700.6052
Internet email at customerservice@ddn.com

Our Corporate Headquarters is located at:
9351 Deering Avenue,
Chatsworth,
California 91311

1.800.TERABYTE, 1.800.837.2298

Product Shipping Instructions

If you are shipping the product to another location, always use the original packaging provided with your unit(s).

If you are sending a product to DataDirect Networks for warranty or out of warranty repair, you ***must*** obtain a Return of Materials Authorization (RMA) number from DataDirect Networks Technical Support.

DataDirect™ NETWORKS

9351 Deering Avenue, Chatsworth, CA 91311

www.ddn.com, 1.800.TERABYTE

© 2011 DataDirect Networks, Inc. All rights reserved.
DataDirect Network, the DataDirect Networks logo, Silicon Storage Architecture, S2A, xSTREAMScaler, and D-MAID are trademarks of DataDirect Networks.
All other trademarks are the property of their respective owners.